

Digital, Culture, Media and Sport Committee

Oral evidence: Fake News – 8 February 2018
(George Washington University, Washington DC),
HC 363

Thursday 8 February 2018

Ordered by the House of Commons to be published on 8 February 2018.

Watch the meeting

Members present: Damian Collins (Chair); Julie Elliott; Paul Farrelly; Simon Hart; Julian Knight; Ian C. Lucas; Christian Matheson; Brendan O'Hara; Rebecca Pow; Jo Stevens; Giles Watling.

Questions 273-620

Witnesses

[I](#): Juniper Downs, Global Head of Public Policy, YouTube; and Richard Gingras, Vice President of News, Google.

[II](#): Monika Bickert, Head of Global Policy Management, Facebook; and Simon Milner, Policy Director UK, Middle East and Africa, Facebook.

[III](#): Carlos Monje, Director, Public Policy and Philanthropy, US and Canada, Twitter; and Nick Pickles, Head of Public Policy and Philanthropy, UK, Twitter.

[IV](#): David Carroll, Associate Professor of Media Design, The New School; Amy Mitchell, Director of Journalism Research, Pew Research Centre; Frank Sesno, Director, School of Media and Public Affairs, George Washington University; and Claire Wardle, Research Fellow, Shorenstein Centre on Media, Politics and Public Policy.

[V](#): David Chavern, President and CEO, News Media Alliance; Major Garrett, Chief White House Correspondent, CBS News; Tony Maddox, Executive VP and MD, CNN International; and Kinsey Wilson, Special Advisor to the President/CEO of the New York Times.

Written evidence from witnesses:

- [David Chavern, President and CEO, News Media Alliance](#)
- [Facebook](#)
- [Major Garrett, Chief White House Correspondent, CBS News](#)
- [Google](#)



HOUSE OF COMMONS

- [Google supplementary](#)
- [Amy Mitchell, Director of Journalism Research, Pew Research Center](#)
- [Dr. Claire Wardle, Research Fellow, Harvard Kennedy School](#)

Examination of witnesses

Witnesses: Juniper Downs, Global Head of Public Policy, YouTube; and Richard Gingras, Vice President of News, Google.

Q273 Chair: Welcome, everyone, to this special hearing of the House of Commons Digital, Culture, Media and Sport Committee. Today's evidence sessions here at George Washington University form part of our inquiry into fake news and disinformation. We are thrilled to be here and absolutely delighted with the support and welcome we have received from George Washington University. We are grateful for their efforts in helping us put on these hearings. This is the first time a Select Committee has taken live evidence in this way outside the United Kingdom, so I am grateful for the work of everyone who has made the logistics behind that possible. I thank the witnesses from the different companies we will see today for their evidence, too.

I have one small piece of housekeeping for people who are unfamiliar with the work of House of Commons Select Committees but perhaps more familiar with the work of congressional committees. It is understood that anyone taking part in a proceeding of Parliament answers questions honestly and truthfully to the Committee, and it is an offence to mislead Parliament. I am not suggesting for a moment that our witnesses would do that, but we do not require witnesses to swear an oath in front of the Committee because, in our rules, it is implied that they will be telling the truth. We do not need to see an oath being sworn to affirm that.

We have a number of sessions to get through today, and I am conscious that we need to do that in good time and in good order. I ask members of the Committee to be clear in directing their questions to individual witnesses and to be clear what they are asking for. Let me also say that, because of the time constraints, I would be grateful if witnesses answered the questions put to them and avoided, as far as possible, being drawn into general statements of policy. If a question is directed to you personally, you should be the person who answers it. I hope that if we can do that, we will keep to some sort of time.

I would like to start with you, Mr Gingras. Do you, and does Google, regard campaigns of disinformation and fake news as harmful to your customers?

Richard Gingras: First of all, thank you very much, Mr Collins, for the opportunity to be here to address the Committee and present evidence.

There is no question that the issue of misinformation and untoward influence on our societies is indeed an issue. It is certainly an issue for Google. I should point out that, as you know, our mission is to organise the world's information and make it accessible and make it useful. Billions of users take advantage of our services every day and rely on us to do the right thing and to provide the right information. So this is not only



HOUSE OF COMMONS

important to us as a company, and important to our users and to society, it is crucial to what we do.

I often think of ourselves as being in the trust business. We will continue to retain the loyal usage of people around the world to the extent that we retain their trust with our work every day. So I agree it is an important and significant issue. We will continue our efforts to make sure that, throughout what we do on Google Search or Google News, we are surfacing the best information from the best possible sources and doing our best to make sure that untrustworthy information does not surface.

Q274 Chair: Thank you. I think we would certainly agree it is important and I am glad that you think it is important too. I asked you whether you thought it was harmful—whether disinformation can be harmful to society and to your customers, who may be innocent recipients of it.

Richard Gingras: Without question, it can be harmful. I think of misinformation in a broad range of areas, not just in news and politics. People come to Google every day looking for, for instance, medical information. It is not hard in the medical area if you are concerned about health. You could take, “Do peach pits cure cancer?”—you can go to Google and indeed you might find information that suggests that. Across our surfaces, as we look at content, it is how we make sure that they have the right information that is helpful, and not harmful, to them.

Q275 Chair: If disinformation is harmful to society and individuals, what responsibility do you think companies like yours and the other tech companies that we will hear evidence from today have? In what sense do you think you have a responsibility to protect your customers from being unwittingly exposed to harmful disinformation?

Richard Gingras: We feel an extraordinary sense of responsibility. As I was noting, the loyalty of our users is based on their continued trust in us, to the extent that if they don’t trust us, they will stop using our products and our business will collapse. So from that perspective alone, it is crucial to us—but it goes beyond that.

We believe strongly in having an effective democracy. We believe strongly in both supporting free expression and supporting a sustainable, high-quality, journalism ecosystem to make sure that quality information is out there and to make sure that our citizens, not just our customers, have the knowledge and information they need to be good citizens.

Q276 Chair: This is to Juniper Downs. Research has been done looking at the role that the “next up” feature on YouTube plays in supplying fake information to users through recommendation. What steps are you taking to address that problem?

Juniper Downs: Thank you, and thank you to the Committee. Good morning, everyone. I am Juniper Downs, global director of public policy and government relations for YouTube.

Our recommendation engine was designed to give users the kinds of content they want to see and it actually works quite well in the main use

cases of YouTube—educational content, comedy, music—providing users with more of the kinds of content they want to see. It is quite useful in those genres.

News content makes up a fairly small percentage of YouTube's overall watch-time—less than 2%. We recognise that there is work to do on our recommendation engine in terms of making sure that we are surfacing the right news content to our users. We have been investing a lot of resources in making sure we surface authoritative news sources and demote low-quality content, but as some of the press coverage has shown over the past week, we still have work to do. There is progress to be made. We recognise that and take responsibility for it.

Q277 Chair: YouTube has supplied the Senate Intelligence Committee with information related to 18 YouTube channels that were backed by the Internet Research Agency in St Petersburg, which account for 1,100 videos and 43 hours of content. Could you confirm whether the identification of those channels was based on analysis conducted by the Intelligence Committee itself, based on information it received, or was that YouTube's own research looking for channels like that?

Juniper Downs: The security and integrity of our products is core to our business model. We co-operated with the congressional investigation into whether there was any interference in the US election, and the channels we discovered on YouTube that were connected to the internet research agency were due to a thorough investigation that we conducted using our own resources. We have publicly reported that information to Congress through the Intelligence Committee, where our general counsel testified back in the fall.

Q278 Chair: Was the identification of those accounts based on your own research, and not on intelligence or information you were given?

Juniper Downs: Our research and our investigation included leads and intelligence provided to us by others, including other companies who were conducting similar investigations.

Q279 Chair: But what did you do yourselves that was based not on leads, but on your ability to analyse how people are using the platform?

Juniper Downs: Correct. We looked at the leads we were provided and went far beyond that in looking at any advertisements that had a connection to Russia. We also looked at organic content to see if there were channels on YouTube that were connected to the internet research agency that were not purchasing advertising but were simply uploading content.

Q280 Chair: You may be aware that the Committee has written to Twitter and Facebook asking for their analysis of whether Russian-backed agencies were involved in interfering with elections in the UK, particularly the Brexit referendum and our recent general election. Facebook is conducting that research for us, and I hope we will get an update from Twitter later, but I don't really know what they are doing about it. Would



HOUSE OF COMMONS

YouTube be prepared to conduct the same analysis for the UK that you have conducted in the United States and look at potential Russian interference in and around the Brexit referendum and other elections?

Juniper Downs: Absolutely. We are happy to co-operate with the UK Government's investigation into whether there was any interference in elections in the UK. We have conducted a thorough investigation on the Brexit referendum and found no evidence of interference. Again, we looked at all advertisements with any connection to Russia, and we found no evidence of our services being used to interfere in the Brexit referendum. We are happy to co-operate with any further efforts.

Q281 **Chair:** Obviously, that is the investigation by the Electoral Commission, which is an agency of the Government. I am asking about this Committee's investigation. We are not necessarily just looking for paid-for advertising linked to the election, but the operation of channels or the uploading of films that can be linked back to Russian agencies and had a political purpose or message during the referendum. Would you be able to do that for us?

Juniper Downs: We are happy to co-operate with that investigation.

Q282 **Chair:** So that's a yes?

Juniper Downs: Yes.

Chair: Good, thank you.

Q283 **Jo Stevens:** I would like to ask about the search function on Google. You may be aware of a British journalist, Carole Cadwalladr. We met some members of your company earlier this week, and she was talking about the auto-fill or auto-finish search. She conducted a simple exercise: when she typed "Jew" into the search function, it would auto-finish the rest of the search for her, and all of it came up with far right, extremist anti-Semitic searches. I am interested to know why it took a UK journalist doing such a simple exercise to identify this issue. Why had Google not identified it previously?

Richard Gingras: I thank you for that question. These occurrences happen from time to time. Many we catch ourselves; others we do not. The auto-complete feature on Google search is an important service to our users: you come to search, you begin to enter a query and we offer suggestions based on what other users have searched for. It is a dynamic feature. Everything we work on in terms of the corpus of expression of the web or the corpus of queries that people ask is a live and vibrant corpus that changes every day. There are indeed also malicious actors out there who will seek to game that as well.

As we have constructed those, we continue to build defences against that, and also to build mechanisms such that users inside or outside the company have the agency to identify these things, mention them and bring them to our attention so that we can quickly correct them. That is what we look to do. Clearly, the kinds of terms that you might see in those

are offensive—offensive to all of us, egregiously offensive—and we will look to take care of them. Again, it goes back to the trust of our users.

Q284 **Jo Stevens:** Your algorithm development and the data they gather is a cyclical thing, isn't it?

Richard Gingras: Yes.

Q285 **Jo Stevens:** So that then determines what phrases are picked up. What safeguards are you putting in place to prevent that cyclical enhancement of hate content, if people are searching using those terms? Can you give us some examples?

Richard Gingras: It is about the continuous advance of our systems, in terms of what kinds of terms and strings of words to look for and making sure they don't occur. It is also about maintaining constant efforts on our part to evaluate the results we surface. We have a large team of what we call raters—10,000 people around the globe who are constantly at work assessing our results against different queries. Are we surfacing the right kind of results? Are we surfacing results that have appropriate authority? They use those hundreds of thousands of bits of data from those evaluations to continue to train our systems. This is an ongoing effort. As much as I would like to believe that our algorithms can be perfect, I don't expect they ever will be, simply because of the dynamics of the ecosystem upon which we work. We will continue to strive to make sure situations like that occur as infrequently as possible.

Q286 **Jo Stevens:** This is a final question on developers and algorithms. Do you have an ethics policy or framework that your developers work within, or do developers' inherent biases—we all have inherent biases—influence the algorithm they build?

Richard Gingras: Absolutely, that is crucial. There are ethics policies that guide our work. For the raters specifically, we have a 160-page set of rater guidelines—it is public information and is available for anyone here to look at—that guides their assessment. We work on top of policy, we constantly evolve that policy and we constantly train those who have to apply those policies to make sure we are doing the right thing.

Q287 **Jo Stevens:** Sorry, can I clarify that? Your raters, as you call them—are they developers?

Richard Gingras: They are not necessarily developers.

Q288 **Jo Stevens:** Right. So do the developers who build the algorithms have an ethics policy?

Richard Gingras: Yes, we have ethics policies. For instance, we have an internal policy called the "honest results policy", which prevents our engineering teams and people like me from trying to influence the algorithms in untoward ways or allow third parties to influence our algorithms in inappropriate ways.

Q289 **Ian C. Lucas:** I was very struck by your statement that you are in the trust business. Trust is based on knowledge. Do you believe your



HOUSE OF COMMONS

customers know what information you retain about them and how you use it?

Richard Gingras: I would hope so. We have made great efforts to provide transparency and control to our users about the information we collect as they use our services. We collect that information to make the services better for them, but they can go to a control panel—hundreds of millions of people have done this—to look at what information we are collecting and, for that matter, to change the settings. They can tell us not to collect certain kinds of information. That is always there for them. It is hugely important that we maintain dialogue with our users about the information we collect from them so we can provide better services to them. We never share that information with third parties—we never have and we never will. It is crucially important that we protect it. Again, it is part of the trust relationship that it is important we maintain with our users.

Q290 **Ian C. Lucas:** So you will give me all the information you hold about me if I ask for it?

Richard Gingras: Indeed. You can come to Google and look at the information we have about you in our different services.

Q291 **Ian C. Lucas:** And will you tell me how you use that information? For example—I am a politician—would you tell me how I could use that information for political purposes?

Richard Gingras: Well, we don't use any information for political purposes. We try to be as transparent as we can about how we use that information. I will give you an example. Google search is not personalised, except that we will tune the results based on where you are geographically. If you are looking for a restaurant in London, we will surface restaurants in London. Otherwise, we think it is important not to personalise search results. We expect that what you want to see and find is any information that is out there in the corpus of expression without us trying to guide you in one way or another based on what we think you might be interested in.

Q292 **Ian C. Lucas:** Do you market your capabilities to politicians?

Richard Gingras: We market various services—for instance, our ads—to various folks who want to advertise using our products. To some degree, that does use data to help target information.

Q293 **Ian C. Lucas:** Do you have a specific team that markets advertising to politicians?

Richard Gingras: I am not that familiar with our sales organisation selling ads to know how we approach these things. We have advertising sales teams in countries around the world. I imagine some focus on different categories, but it is not an area of my expertise in my role as vice president of news.

Q294 **Ian C. Lucas:** Is it within your area of expertise, Ms Downs?

Juniper Downs: I want to reiterate some of what Richard said about how we use the information we collect from users. The main principles of privacy at Google are transparency and control. For YouTube, for example, we do collect the watch history of signed-in users, but the user can at any time pause or delete that watch history. The way we use that information is to improve the service for users. So, earlier we were talking about recommendations. If we know that someone is a lover of comedy or a particular kind of music, we can use that watch history to optimise the service. We never sell the information to advertisers.

We do provide aggregate data to advertisers, to help them optimise their campaigns. In terms of the work our sales team does with various constituents who are interested in using Google's ad products, that is the kind of information that would be provided—never information about individual users.

Ian C. Lucas: Thank you.

Q295 **Chair:** On that point, if you can provide such precise aggregate data, why do you find it so difficult to identify bad content on the platform? Why do you still have problems of fake news being filtered in next up?

Juniper Downs: Identifying and managing content on YouTube is the No. 1 priority for us this year. It is mission-critical for the business. It is critical to our users, to our creators, to our advertisers and to us as a company, so we invest tremendous resources in terms of both technology and the people working on these issues. Our executive team is absolutely engaged. We meet for hours every week to figure out how we can improve our systems to make sure that the policies on YouTube are followed and that we quickly identify content that violates them and removing it. This is a top priority for the business.

Q296 **Chair:** What percentage of your advertising revenue do you reinvest in the site in this way?

Juniper Downs: I do not have an exact percentage, but we have spent tens of millions of dollars fighting spam and abuse across our products. We have committed to having 10,000 people across Google by the end of this year working on these issues.

Q297 **Chair:** Tens of millions of dollars?

Juniper Downs: Tens of millions.

Q298 **Chair:** What is the annual advertising revenue for YouTube?

Juniper Downs: I am sorry; I don't know the answer to that question.

Q299 **Chair:** Just roughly—you must have a rough idea.

Juniper Downs: I am sorry; I don't know the answer to that question, but I want to reiterate that we will invest the necessary resources to address these issues not only in terms of people we employ but the technology we develop. We are continually staffing up and we have seen good progress in our management of these issues. For example, over the



HOUSE OF COMMONS

past eight months, I think the work we have done on violent extremism speaks for itself: not only have we removed over 100,000 videos but the speed with which we have been able to identify this content and remove it from the site has gotten faster and faster. We are now at the point where 70% of that content is removed within eight hours of upload and 50% within two hours of upload. That is the kind of progress that we demand of ourselves and that we will continue to strive for across all of these issues as we move forward.

Q300 Chair: We will see if we can work out during the course of the morning what that percentage is—my suspicion is that it might be quite small. Many people still say, though, that you can sell the service to micro-target people—that is how you make your money—so you must be capable of using the same tools to more effectively identify harmful content. Of the content you do take down, though, what proportion of that is user referral of bad content and how much is bad content you discover for yourselves?

Juniper Downs: The technology we have developed to identify content that may violate our policies is doing more and more of the heavy lifting for us. In the area of violent extremism, 98% of the content we have removed has been identified by our algorithms. That number varies from policy to policy, depending on how effective and precise the technology is on particular issues. When it comes to misleading and deceptive content, our spam detection systems are quite effective at identifying the behavioural patterns associated with mass-produced misinformation that is being distributed through our services. In 2017, we removed over 130,000 videos for violating our misleading spam policy. Those were virtually all identified by our technology.

Q301 Chair: If you had evidence of someone being heavily exposed to or distributing violent extremist content or other forms of extreme content that might be dangerous to children, what sort of responsibility do you think you might have to share that knowledge with law enforcement agencies and people who might take an active interest in that person?

Juniper Downs: We co-operate and meet regularly with law enforcement to share information. When it comes to the disclosure of user data, obviously we do that pursuant to valid legal process. There are mechanisms in place where law enforcement can request information from us, pursuant to their investigations. There are also emergency disclosure provisions, where if we identify content on our services that poses an immediate risk to life, we disclose that information proactively.

When it comes to the safety of children, obviously we report to the National Centre for Missing and Exploited Children any solicitation of minors or content that is exploitative in that way. It co-operates with law enforcement internationally to make sure that incidents are followed up on.

Q302 Chair: I appreciate that you co-operate with authorities upon request. If someone had a family member who was a victim of a violent act carried



HOUSE OF COMMONS

out by someone who, from their interests on YouTube—perhaps the channels they have set up—quite clearly was potentially quite violent and dangerous, some people might ask whether you have a responsibility to notify the authorities of people who give you cause for concern, and not just wait to co-operate with an investigation after a crime has been committed?

Juniper Downs: Certainly, if we identify someone on our services who is posing that kind of threat to individuals. We see instances where someone may be threatening to perpetrate a violent attack, we would proactively disclose that because we would consider that part of the emergency disclosure provisions.

Q303 **Rebecca Pow:** I want to look at the ethics of YouTube, your “up next” button and your scoring system for videos. Would it be right to say that, usually, the content with the highest score is something that has had the most views? Are hits important?

Juniper Downs: The recommendation engine is quite complex and it varies depending on the video that is being watched. One of the factors is content that is popular on YouTube. There is also content that is associated, so if I have a real niche interest in a particular kind of knitting, there may not be a lot of highly popular videos that are about that type of knitting, but we will continue to recommend videos that are similar and that provide more instruction on that type of knitting.

Q304 **Rebecca Pow:** Popular usually means lots of hits, doesn’t it? I want to mention some research that was done recently by *The Guardian* that suggested that the largest source of news for false news channels came from these videos on YouTube that tend to be recommended by the “up next” button. More research by Oxford academics demonstrated that a lot of that had a far right bias. I put it to you that if that is the case, you are indoctrinating society.

Juniper Downs: I am not here to comment on the particular methodology of *The Guardian* research, but when it comes to recommendations, we recognise that we have more work to do. Recommendations are largely a reflection of user interest. When it comes to news content, our goal is to promote more authoritative sources and to demote lower-quality content from less established sources.

It is not only popularity that fuels the recommendation engine. It also tries to draw from well-established, known sources of news and put that near the top. Do we succeed at every time? We do not. I will give you a recent example. Earlier this week when the stock market fell, our algorithms recognised “stock market plunge” as being a newsy query because established news sources were using the word “plunge” in their coverage, but they did not recognise “stock market crash” because, for some reason, that term wasn’t being used. That was flagged to us and we immediately corrected it and made sure that we were surfacing authoritative sources across both of those queries, but even those subtleties in language can cause our systems to misfire on occasion and we work quickly to try to correct that.



HOUSE OF COMMONS

Q305 **Rebecca Pow:** *The Guardian* information fed 8,000 YouTube recommendations into their software system, which tracks political disinformation campaigns. In every case, the largest source of traffic for these fake news channels came from YouTube's "up next" algorithm. Are you clearly saying that you do not agree with that?

Juniper Downs: I am saying it is an area of progress for us. That research was done in the lead-up to the election. We have made a lot of progress since in making sure that the more authoritative sources surface near the top of our recommendation engine. It is an area that we will continue to work on and improve. We recognise that some of those results are not where we want to be as a product, so it is an area of investment for us.

Q306 **Rebecca Pow:** There is one other area I want to touch on. It seems like, in choosing to rank videos and to have a system of what is good, what is better, what is higher and what is lower, you are in fact acting as an editor—that is what an editor of a newspaper would do—yet you are not calling yourselves editors; you are calling yourselves hosts. I wonder whether you think the description of what you are and the whole name of your platform ought to be changed so that you take on more of the responsibilities of a bona fide newspaper and you have to apply broadcast and newspaper regulations to yourself. At the moment, you are unregulated.

Juniper Downs: When it comes to the scale of our services, we have well over 400 hours of content uploaded by creators to YouTube every minute. There has to be some organising principle. When someone comes to YouTube and searches for "funny videos", there has to be an algorithm that delivers the most relevant and useful results to that individual user. That is how our algorithms function. The question of whether we are a publisher really gets to the question: do we take responsibility for the content on our site? The answer to that question is absolutely yes. It is a top priority for us to make sure that we are providing a useful service to our users and that we are exercising responsibility.

Q307 **Rebecca Pow:** But you are self-policing, aren't you? Is that right?

Juniper Downs: We have a set of community guidelines that set the rules of the road for YouTube. Those have been in place since we acquired YouTube as a service. They are dynamic. They evolve in response to changing trends in the world and changing trends in what kinds of content and misuse of our services we see. We have developed those guidelines because we are trying to maintain a certain kind of community for our users. They do go above and beyond what the law requires, but we also follow the law in the jurisdictions where we are launched. When we are alerted to content that may violate a particular law, we block that content for the relevant jurisdiction.

On the model of self-regulation I have just described, the EU Internet Forum, where we have been co-operating with the European Commission for several years on the removal of illegal hate speech from our services, is a good model for how Governments, NGOs and the tech industry can



HOUSE OF COMMONS

collaborate to ensure that we act responsibly and quickly to deal with these content issues. The Commission has acknowledged the progress of that process.

Q308 Rebecca Pow: You are described as “the most overlooked story of 2016”, and people are describing this whole social media issue as a disease on society that is affecting behaviour. It is surprising that you do not seem to register any of that.

Juniper Downs: We take the criticism there very seriously. Again, it is a top priority for the company to make sure that we continue to invest in detecting content that may violate our policies and watch for bad actors who try to manipulate or exploit our systems. The openness of YouTube has brought tremendous benefits to society when it comes to educational content, culture, music and the arts. Those benefits are tremendous. Of course, that openness also presents challenges, and those challenges are ones that we are committed to managing and dealing with responsibly.

Q309 Chair: The Committee Clerk has been busy trying to work out the answer to the question about how much money YouTube makes. Independent market analysis—I appreciate the figure is not disclosed by the company—estimates that YouTube’s global revenues are about \$10 billion. Does that sound about right?

Juniper Downs: I really cannot confirm any of those numbers. It is not information that I have access to.

Q310 Chair: Well, let us say it is, for easy maths. If you are investing \$10 million a year in making YouTube safer and your revenues are \$10 billion, you are reinvesting 0.1% in making the site safer. If that is true, it sounds to me like a small sticking plaster over a gaping wound, and it is not likely to be satisfactory in addressing some of the concerns that Rebecca Pow has raised. I do not know whether you want to comment on those figures. I appreciate the company may not want to disclose revenue figures, but if it could give us a percentage figure for the amount that is reinvested in the site as a proportion of revenues from YouTube, we would be very grateful for that. As I say, based on publicly available figures, 0.1% seems to be where that comes out.

Juniper Downs: I think your question gets to whether we are investing adequate resources in addressing these issues. There really is no constraint on the resources that we will put towards getting this right. We have invested tremendous time—time from our engineering teams, from our executive teams and from the trust and safety operation, which has nearly doubled in size over the past year. We take these issues very seriously. If we decide that we are unable to make our desired progress because we do not have adequate investment, we will invest more in addressing it.

Chair: All I would say is that, for a multi-billion-dollar business, if tens of millions of dollars is the investment to address these widely held and increasingly well-documented concerns, it sounds like a pretty unambitious programme of investment. I think, understandably, given the

huge size of the platform, people would question whether that sort of investment is ever likely to be successful in addressing some of these concerns.

Q311 **Christian Matheson:** Ms Downs, I want to follow on from Ms Pow's questions about the Paul Lewis article in *The Guardian*. It quotes Guillaume Chaslot, one of your former engineers, who says that "Watch time was the priority" for your algorithms in prioritising what came up next, not necessarily truthfulness, decency or honesty. How would you comment on that allegation?

Juniper Downs: Watch time is obviously an important metric for us, because it demonstrates that we are providing a service that users love, where they want to spend time on the product or they are enjoying the experience of YouTube and finding it valuable. For certain kinds of content, such as watch time for music, I know I listen to hours of music on YouTube because it keeps delivering the music I want to hear. We know that for news and for certain other verticals, the veracity of the content we provide to our users is also incredibly important.

That is why some of the changes we have made over the past years have been about making sure we are surfacing more authoritative content to users, where they are getting the information they are seeking, and demoting lower-quality content. We have created dedicated new surfaces. We have a new breaking news shelf, so when there is a breaking news event we surface high-quality content near the top, and then the recommendations are also limited to those authoritative sources.

We also invest in making sure that the users of YouTube have the skills they need to assess the content they are watching. In the UK, we created a campaign called Internet Citizens, where we are committed to training 20,000 young people in identifying fake news and learning how to check sources and bias and so on. It is about an effort that supports the whole ecosystem.

We also work with publishers. We launched a programme called Player for Publishers, which now has 50 European news providers in it. We provide them with the technical support to have a video player embedded in their own website, which reduces complexity and cost for them, and we also work with them to optimise the content they produce for YouTube. Their watch time has doubled, and they really appreciate the fact that their watch time has increased through our investment in them. Watch time can be an important metric.

Q312 **Christian Matheson:** So you do accept that there is a problem of bogus news, misleading news and disinformation, and you also accept that you have both a responsibility and the means to address that?

Juniper Downs: We recognise that there is a problem of misinformation and we are dedicated to making sure that we are promoting authoritative content when it comes to news and demoting misinformation and low-quality content.

Q313 Christian Matheson: Do you have a sense of moral responsibility to take those actions?

Juniper Downs: We have a sense of social responsibility and it is a business priority for us. The trust that users have in our services is reliant on our providing a high-quality experience when they come to the site. Presumably, when people come looking for news rather than entertainment, what they expect to find is factual reporting from a variety of sources. We do provide information from diverse news sources from around the world. We think that is valuable. So yes, our investment in this area is because we recognise that it is an important policy matter, but also because it is an important business priority.

Q314 Christian Matheson: Paul Lewis confirmed the Chaslot conclusions in relation to the United States presidential election, which suggested that YouTube was six times more likely to recommend videos that aided the Trump campaign than the Clinton campaign. Do you accept those figures?

Juniper Downs: Again, I am not here to comment on the particular methodology, but I will say—

Q315 Christian Matheson: No, go on, comment on it.

Juniper Downs: We did publicly say that we did not agree with the methodology that was used.

Q316 Christian Matheson: Why not?

Juniper Downs: We were not provided the 8,000 videos to examine them, but this is an employee who left the company—

Q317 Christian Matheson: You've got a bigger selection, because you have got the whole set, so maybe you can run their methodology past your bigger set and get what you might consider to be a more accurate solution. How about doing that?

Juniper Downs: Let me explain how I think the recommendation engine functioned in the lead-up to the election. It is a reflection of user interest, so to the extent there was more content recommended about one candidate versus the other, that was because there was a lot more user interest being expressed about that candidate. So we saw more searches for that candidate.

I think we saw similar trends across broadcast coverage of the US election, where one candidate got more coverage than the other. So that is a reflection of what users wanted to see, what they were interested in and what they were looking for when they came to the site.

Q318 Christian Matheson: But Ms Stevens has already identified the possibility that that becomes a self-fulfilling prophecy, because your recommendations and "up next" simply reinforce what has gone before, and it tends to spiral in one direction, as opposed to providing balance.



HOUSE OF COMMONS

Juniper Downs: This is an important question for us to figure out how to ensure that we are providing information that users actually want to watch and will watch—that is also diverse. We do not build any political bias into our algorithms. We design products for everyone and we invest heavily in ensuring that there is no bias built into the algorithms we design.

There is a human element here, too, though. People want to watch what they want to watch, so if someone is not expressing an interest in a particular person or type of content, it is hard to insert something opposite to that and expect them to watch it. We see abandonment of the service when we do that, because humans have their own will. So we try to create a diversity of content but, at the same time, keep it topical enough that the user feels we are fulfilling the interest they have expressed to us when they come to the service.

- Q319 **Christian Matheson:** Finally, if a video is getting hundreds of thousands or even millions of hits and, all of a sudden, when it is at the top of its performance, it gets taken down, would that not be a little strange, in reference to your business model.

Juniper Downs: I am not sure I quite understand the question.

- Q320 **Christian Matheson:** If a video is getting hundreds of thousands or millions of hits—which is good for you and for the person who posted the video—and then they take it down when it is at the top of the game, would that not strike you as a little bit odd?

Juniper Downs: The community guidelines and content policies we have for YouTube apply to all of our creators equally, no matter how popular they are. We identify content when it is flagged by users or detected by our own internal systems. We evaluate those videos and, if they violate our policies, we remove them immediately. It does not matter who the creator is or how many views the video has.

- Q321 **Christian Matheson:** What if the creator takes it down?

Juniper Downs: The creators on YouTube always have control over their own content and are free to delete it at any time.

Chair: Thank you.

- Q322 **Julian Knight:** You have talked about how much progress you have made, and you believe that your moves in this area are quite effective, yet we have just discovered that potentially you are spending 0.1% of your turnover. More recently, in fact this morning, the front page of *The Wall Street Journal* has a forensic investigation—old-style journalism, if you like—by Jack Nicas, who says that your recommendations “often present divisive, misleading or false content, despite changes the site has recently made to highlight more-neutral fare”. Potentially up to 70% of viewing time is taken up with these recommendations. In the light of that, why has your self-regulation so demonstrably failed, and how many chances do you need?



Juniper Downs: Again, our recommendation engine was designed with the main use case of YouTube in mind—the things that people come to YouTube and love the site for: comedy, music, cooking vlogs and beauty vlogs. It works quite well for those use cases. We recognise that for news the effort we have put in to promote more authoritative sources and demote lower-quality content is a work in progress. We will continue to invest in getting it better.

Q323 **Julian Knight:** But you stand by the fact that it is quite effective? That is what you said at the start.

Juniper Downs: It is quite effective for the majority use case of YouTube. But news is less than 2% of our watch time, so it is actually not the majority use case of the platform, but is an area—

Q324 **Julian Knight:** So you have got the 98%, but you think for that 2% your algorithms do not work at the moment.

Juniper Downs: I think they are working better than they were six, eight months ago and so on. This is an area where we are continually investing to make sure that we are providing the right news experience for users when they come to YouTube. We are committed to doing better. When we see some of the results in that *Wall Street Journal* article, frankly we are not proud of them. That is not the experience that we want to provide to our users—

Q325 **Julian Knight:** So you agree with the methodology of this article and take that on board? Potentially, you dispute the one from the past in *The Guardian*. We have a phrase in Britain—"It's a fair cop." Is this a fair cop?

Juniper Downs: I think the results pictured in that article were taken from YouTube, so to the extent that we look at those and think we could do better, we certainly look at those and think we can do better.

Q326 **Julian Knight:** One area where you seem to do remarkably better is over things like sporting rights. For example, if I wish to post a final touchdown from the Super Bowl—I am saying that as I am in America, but I am probably thinking about a premier league soccer goal—that would be taken down within minutes. It would be almost instantaneous, and you would be right on to it. Why does it take so much longer, if at all, when it comes to misinformation by foreign powers that are specifically looking to undermine the west?

Juniper Downs: When it comes to copyrighted material, we have invested a lot over the years to make sure that we protect people's intellectual property rights. Part of why the detection and removal is so much faster is that rights holders provide us with a digital file of their copyrighted material, so we have that as the starting place. We already know what content the systems are looking for, so we can identify matches with that content relatively quickly using our technology.

If we were provided with a digital file of every misinformation video, we would be able to identify them just as quickly. Unfortunately, our systems have to look for patterns, and they are quite sophisticated at doing that,



HOUSE OF COMMONS

but bad actors are also constantly evolving and trying to evade detection. It is a little bit of a cat and mouse game, and we are obviously always trying to stay one step ahead of those who attempt to misuse our services.

Q327 Julian Knight: So it is not just about the money or the fear of being sued that you act so quickly when it comes to copyright? However, when it comes to misinformation from, as I say, foreign actors, the evidence is abundant and out there that it is effectively not in the same ballpark, to stretch the sporting analogy.

Juniper Downs: It is a top priority for us. As I said, the security—

Julian Knight: 0.1%.

Juniper Downs: I am not sure that the financials behind this are the right metric, in that the technology we deploy, which is a hard thing to monetise, actually does a lot of the work at scale to identify this content with speed. If you try to quantify in people hours how many hours of human endeavour are being saved by the technology, we would get to a much vaster number. In order to address these issues at scale, we really need to invest in the technology that can find the content and enable us to act on it quickly.

Q328 Julian Knight: Okay. I will turn to you for a question in a moment, Mr Gingras—I am sorry to leave you out. Just ahead of their elections, Germany effectively regulates in terms of removing hate speech from social media. I think they were concerned about the reaction to immigration. Social media responded, and many commentators suggested that there was a demonstrable reduction in the level and effectiveness of their elections—it also happened in the French elections. Surely that is strong evidence that the way in which western democracies protect themselves is to regulate you?

Juniper Downs: We work with elections and campaigns around the world to help them with digital tools to protect their websites and campaigns from any interference—from phishing attempts and so on. We are very committed to investing in the integrity of elections and to providing security support to political candidates and campaigns. That is a service that we can help to provide. We collaborate with others on that.

In terms of the German law that was recently passed, there has been robust public debate about the risks and benefits of that approach. I think, since it was implemented in January, there has been a lot of public criticism of the way it has played out. The debate is ongoing on whether that is the right approach.

Q329 Julian Knight: Mr Gingras, your company very kindly hosted us in New York in order to talk about the Trust Project, which was very interesting. While we were in New York, we also received some quite stark evidence of the crushing impact on local press, and increasingly on national press, of the sucking out of advertising revenue by companies such as yours from more traditional media.



HOUSE OF COMMONS

You talk about how you want to train journalists so they are able to use your platforms more effectively, but what happens when all we have are a few global players, such as *The New York Times*, *Washington Post* and, on our side of the Atlantic, the British Broadcasting Corporation? It is pointless having a journalistic programme if there are so few journalists left. What else will you do in order to ensure that proper journalism survives?

Richard Gingras: That's a very important question, and it is one we are deeply concerned about as well. I should point out that the success of our efforts with Google Search and our ads business is really dependent on there being a rich ecosystem of knowledge out there, including at the local level. As we often say, our success is dependent on the success of publishers. I will point out that the publishing business, in every dimension, has without question been changed and affected by the introduction of the internet. I will also point out, however, the various things that we are doing to assist in this.

First, we share 70% plus of our display ad revenue with publishers around the world, to the tune of some \$12 billion—that is one element—as well as providing them with ad systems and tools that allow them to derive revenue. Some 2 million publishers around the world benefit from that. We also drive traffic to many new sites, to the tune of well beyond 10 billion visits per month, which has been valued by third parties at between 5p and 7p per visit. That is another \$5 billion to \$7 billion a year in revenue to the industry. So that is one component, but it is just one.

We need to recognise that the ecosystem has changed. One way to think about this is through our own actions. I started out in the newspaper business—I started out in the press rooms of *The Providence Journal* when I was a teenager.

Q330 **Julian Knight:** Does *The Providence Journal* still exist?

Richard Gingras: Yes, it does still exist.

Q331 **Julian Knight:** There are not that many local newspapers in the US that do. It is good that you are giving this background, but the point is what happens when the journalistic landscape is so decimated and people do not know what to trust? Do you have a sense of responsibility to ensure that we have that trusted news?

Richard Gingras: We do feel a strong sense of responsibility, for the reasons I pointed out. It is important to society and it is important to the nature of our business, so we have mounted many efforts. My role at Google is, on the one side, to oversee our efforts to surface news on Google Search and Google News, but another part of it is the various efforts we have mounted to enable the publishing systems and legacy publishers to make the transition from the old marketplace of information before the internet to today's one.

Behaviours have changed. Think about it. When I was young and looking for a used car, I went to the classifieds in my local newspaper. If you are



HOUSE OF COMMONS

doing that in Birmingham today you will go to Gumtree, not your local newspaper. If I was looking for a house or a rental apartment as a kid, I would go to a newspaper. Today, in Liverpool, you will go to Zillow. That is even when it comes to national news. I went to local newspapers for my national news—much of it might have been from wire services, but I went to them. Today, people's behaviours have changed. They go to *The Telegraph* or the *London Times*. The marketplace has changed.

We feel that part of our responsibility is how we help and enable news organisations to figure out what kinds of products they can serve to their local communities. We do that in various ways. One is through deep collaboration with the publishing industry. It is important that we understand their challenges and they understand what we can do. We have mounted efforts such as the subscription project. Can we bring better tools to help to drive subscription revenue? Can we bring better data and knowledge to help them to identify targets of opportunity in their markets? We helped with tools for digital storytelling and training for journalists through the Google News Lab. We have trained thousands of journalists in the UK, and will continue to do more. We have provided innovation funding in Europe through the digital news initiative—some \$83 million to date, and £6 million, I should say, in the UK, helping folks like Trinity Mirror develop news services like InYourArea.

There is a lot of opportunity to continue to develop, and we are seeing good signs of success around the world, but there is no question that there is more work to be done and more innovation that needs to take place to reach that sustainable ecosystem of local news publishers as well.

Julian Knight: Thank you.

Q332 **Brendan O'Hara:** I have a couple of follow-up questions—one on Mr Matheson's and the other on Mr Knight's. Following on from Mr Matheson, your business relies on people being on your site and viewing for as long as possible. Hence, you have these sophisticated recommendation engines, which are designed to do that. How do those recommendation engines differentiate between children and adults? Is it the goal to get the user, regardless of age, to stay on the site and consuming content for as long as possible?

Juniper Downs: When it comes to Google, we only allow users who are 13 and older. That is the limitation on the site. Your question gets to the issue of public concern around tech addiction, and particularly young people's use of smartphones and social media. When I think about the goal of Google Services, it is really to provide products that enhance people's lives, not detract from them.

We think the question of how we can better fulfil that goal is a really important question for society, so we are investing in research to better understand these issues and to better understand what kinds of product design we can implement to make sure that we are providing a product that is an enhancement to people's lives and enhances people's productivity. If you look at YouTube, we have 1 billion views of educational

content every day. That is a tremendous asset to society. If you think about something like Khan Academy, where an uncle started making videos—

Chair: Sorry to interrupt, but I need to direct you back to the question and away from making general statements about YouTube.

Q333 **Brendan O'Hara:** Thank you, Chair. The question is how you differentiate between adult consumption and child consumption.

Juniper Downs: We don't look at adult consumption versus child consumption. The recommendations are based on the video that is being watched and content associated with that video, or with the watch history of the individual signed-in user. Again, only people aged 13 and older can sign in to YouTube, so the watch history there would be teenagers and above.

Q334 **Brendan O'Hara:** Okay, so there is no mechanism, other than an arbitrary "You have to be 13 to view", that allows you to say who is watching what? Anyone of any age—any tech-savvy five or six-year-old—could probably get around that 13-year-old barrier. Any tech-savvy child could consume as much content as they like and you wouldn't know.

Juniper Downs: Protecting teen users on YouTube is a top priority for us. One thing we do is age-gate videos that are more mature. You have to be signed in and aged over 18 to view those. That is the one place where more mature videos aren't eligible to be viewed or recommended.

Q335 **Brendan O'Hara:** I am not talking about more mature videos. I am talking about the unrestricted consumption of YouTube videos by children. Is there a way that you have managed to address that? I suspect from your answer that you haven't.

Juniper Downs: The investment that we have made is in developing a dedicated YouTube experience for families and children, which is a dedicated app called YouTube Kids. When we designed that product, the goal was to create a new service that gives parents control and information. It was informed by collaboration with child safety experts.

We designed a product that is a much more limited corpus of YouTube videos—the family-friendly content—and we give parents controls such as a timer for how long their child can view the app, the ability to turn search off, so that the child is limited to a much smaller set of videos, and other controls. That was because we want to provide a good experience for families, but we recognised that it really needed to be in its own space.

Q336 **Brendan O'Hara:** Are you happy that you are doing everything you possibly can—this goes back to the trust issue that you talked about, Mr Gingras—to protect children from this unlimited consumption that we hear so much about? That is why people have so many issues about child interaction with the internet at the moment.

Richard Gingras: I would point out that we are never fully satisfied with our work. As I have mentioned, this vibrant, live ecosystem changes every



HOUSE OF COMMONS

day. For instance, with Google Search, we get billions of queries every day, and 15% of those are queries we have never seen before. It is an ongoing effort. As I often mention, our algorithms are not perfect and I doubt they ever will be, which means we have to be ever vigilant.

I noted earlier, when we were talking about how many dollars we have invested in this, that when we mention, for instance, \$10 million of additional investment in security efforts and so on, that is just additional recent investment. It is building on a foundation of effort. We have thousands of engineers at YouTube, thousands of engineers at Google Search, security systems and security people that we are building upon, and will continue to build upon, such that we can understand the new challenges and address them as quickly as we can. I just want to be clear that there are no laurels that we rest on here.

Q337 Brendan O'Hara: My final question relates to what Mr Knight was saying. I am talking about the fairy-tale idea of a fun space where everyone could share and enjoy. It really has been superseded by events. It is no longer possible to have YouTube operate in the way it did way back at the beginning, because of people who are exploiting it and social issues that have arisen. Are you satisfied that the current legislative framework is sufficiently robust to meet these new challenges, or would you welcome further legislation; or, as the debate is now happening, how would you like to participate in that debate, to shape a legislative framework that works for everyone?

Juniper Downs: We are intrinsically motivated to address these issues. Again, it is the top priority for YouTube as a company, and we welcome any conversation about legislative proposals. A big part of my organisation's role is interacting with policy makers and other stakeholders to hear ideas and to participate in conversation about them, so we are always happy to do that. I want to reiterate that we do not need extra motivation to get this right. It is mission critical for the business for us to address these issues responsibly, and that is what we are committed to doing.

Richard Gingras: But more collaboration, more engagement—these are problems that we are facing as a society, that we are all facing. There are no silver bullets for any of these questions; they are big questions, whether we are talking about the challenges of misinformation or of the evolution of the publishing ecosystem. If there is anything that we have learned over the last several years—this didn't start in the last 18 months—it is that our best approach to these is to engage with the community, whether they be publishers, people in the public policy space or whoever, such that we can all, collectively, gain a better understanding of how to approach the problems and what they are and how we partner together on finding better solutions.

Chair: Thank you, Mr Gingras. We take your point, but we have heard the expression "top priority" a lot. If we judge the company based on what it does rather than what it says, the top priority is maximising advertising revenue from the platform, and a very small proportion of that is



HOUSE OF COMMONS

reinvested into dealing with some of the more harmful content. That is one of the reasons why we are here and why social concern about this is growing. Giles, do you want to come in?

Q338 Giles Watling: To use pretty much every analogy under the sun, it seems to me that you have opened a Pandora's box, probably unintentionally, and you have got a tiger by the tail. And you are now, to carry on the sporting analogy, somewhat behind the curve.

It seems extraordinary. You were talking about misinformation. Let's face it: misinformation, disinformation, fake news can be sort of sexy. People like ghosts, conspiracy theories and UFO stories, and they follow it. This creates, I would imagine, a huge amount of traffic, both propagating those theories and so on, and from those who want to quash it. That would be good for you, because you can then harvest the data from the people who are doing that, and that would have a commercial and potentially political value, would it not?

Juniper Downs: Let me talk a little bit about the spectrum of fake news, because the term is used to reference a lot of different categories of content, and then I can speak to how we deal with it and address your question.

On one end of the spectrum, you have fake news farms, content that is produced in bulk, which there is a financial incentive to try to distribute as virally as possible. This pattern is caught by our spam detection systems. It violates our policies; we remove it.

Then there is what I think you may be referring to, which is clickbait—content that might be more sensationalistic. We see videos with titles in all caps, or salacious thumbnails. We see this clickbait content. What we do there is demote it. Our algorithms are trained to identify clickbait and demote it. It is curious that you ask the question about watch time and user interest. What we expected, when we started demoting clickbait content, was that that would create a decline in watch time, and that was the initial result. We were willing to absorb that, because we did not think that it was high-quality content, so we were willing to absorb the loss of watch time, but—this may make you feel better about humanity—over time, the watch time actually picked back up and increased. So when we were surfacing the higher-quality content, people actually started consuming that as well.

If you think further down the spectrum of what people call fake news, you also have, somewhere in the middle, various opinions that people express on YouTube about current events, political matters and so on. We do not consider that content to be news. It is subject to our content policies. If it crosses the line to hate speech or harassment, obviously we are going to take it down, but that is not news content, in our view.

Q339 Giles Watling: But you did say a little while ago that people watch what they want to watch. You are sort of on the back of that and you are not able to keep up with it. I put it to you that it is actually not in the interest of large corporations like yours to put a stop to this, because you are

harvesting data.

Juniper Downs: This is really about the kind of site and service we want to provide to our users and the long-term play. It is not necessarily about short-term results. When it comes to news content, our commitment is to provide our users with authoritative information from established news sources. We want to demote that lower-quality content, not necessarily because we think that we will get more watch time. Your assumption that people will watch more clickbaity content because it is entertaining, or for whatever reason, may be accurate. That is not enough for us to prioritise it. Our goal is to demote it and to promote the more authoritative sources because we think it is a better long-term play for the business to be seen as a trusted provider of news and information to our users.

Giles Watling: Thank you.

Q340 **Chair:** A final couple of questions from me before we close this session. I notice that YouTube's policy now is to put a label or a tag on content from broadcasters that are publicly funded in their countries, I think in response to requirements in the United States that that be done. Is that correct?

Juniper Downs: We recently started rolling out a new transparency feature for certain news sources. We made a public commitment to provide users with information about the sources of news they consume. We have started fulfilling that commitment by introducing a new label that informs users if a news outlet receives Government funding, and a separate label to inform users if a news outlet is a public broadcaster. It has just started rolling out, but we plan to extend that and think that providing transparency to users is an important part of what we can provide as a service.

Q341 **Chair:** There has been some criticism of this, because it means that RT, a Kremlin-backed propaganda station, is in the same bucket as the BBC, a publicly funded, editorially independent public broadcaster. Do you think mixing RT and the BBC together is making navigation easier for consumers or muddying the waters?

Juniper Downs: As we developed this new feature, we did reach out to various news outlets, including the BBC, and we got that feedback directly from them, which is why we created two separate labels: one that states that an entity receives Government funding in full or in part—that is what applies to RT—and a separate one that labels things as public broadcasters, which is what applies to the BBC. There are actually two distinct labels. Each label links to an online source where users can read more about the news source and get additional information about who funds them and what their structure is.

Q342 **Chair:** There is a degree of confusion here. I think some people reading that might say, "What's all the fuss about RT? They're a public broadcaster. They're state-backed. So is the BBC. It's all the same thing." Actually, they are radically different things. It is not just about the ownership structure or the model of the organisation but about the way it



HOUSE OF COMMONS

is run and the degree of editorial independence it has, which exists at the BBC and does not exist at RT. Certainly from the evidence we have received from people while we have been in America this week, it is very strongly suggested that this is something that needs to be thought through again, because it may unwittingly sow more confusion rather than adding the transparency that you have asked for.

Juniper Downs: We agree that more information and context is better, which is why the label links to a site where people can get additional information, like what you describe, to distinguish some of these news providers from each other.

Chair: That's all well and good, but given the way people consume content—we heard earlier about "Up Next" providing a constant stream of content for people to see—I am not sure that that is a service that people in reality will use that much of. What they will go on is the labels. If labels seem to be categorising content in a similar way, they may trust them in a similar way, and that may not be helpful to them in terms of making sure they are consuming the highest quality news.

I think we have completed our questions for this session. Richard Gingras and Juniper Downs, thank you very much for your evidence.

Examination of witnesses

Witnesses: Monika Bickert, Head of Global Policy Management, Facebook; and Simon Milner, Policy Director UK, Middle East and Africa, Facebook.

Q343 **Chair:** We will go straight into the second panel. Monika Bickert and Simon Milner, thank you very much for joining us here this morning for our evidence session. I will start by asking Monika Bickert for some background information about the Senate Intelligence Committee's investigation into Russian activity during the election. Quite a lot of that work has been quite helpful and insightful for us and, as you will know, it has obviously led us to contact Facebook and ask for a similar analysis of whether there was any Russian activity in the UK.

I want to ask about the information that Facebook has provided. Is it correct to say that all the evidence so far—the analysis of the number of people who were exposed to content created and distributed by the Internet Research Agency in St Petersburg—comes from the analysis that Facebook did, identifying payments made in roubles to promote advertising around that content?

Monika Bickert: Thank you for the question. The information that we gave to the Committee in testimony—it was my colleague Colin Stretch—is really the best source of information for the details about that. In that testimony, Colin described how Facebook conducted that investigation. Ads being paid for in roubles was one of the signals, but we considered other things as well. Mr Stretch's testimony is really the best record of that.

Q344 **Chair:** From the information that we have received from the Senate



HOUSE OF COMMONS

Intelligence Committee, it is very clear that everything has been extrapolated only from those accounts where rouble payments were made. Analysis has been done to try to identify accounts linked to that behaviour, but Facebook simply analysed what some people might call the lowest-hanging fruit—linking it from the accounts where rouble payments had been made—and did no wider analysis of whether other agencies were involved or other similar activity took place that was not linked to rouble payments for promotional advertising.

Monika Bickert: I think I can just refer you to Mr Stretch's statement. We also put out some posts, including in April 2017 and several over the summer—I believe one in August and two in September—in which we talked about how we looked for that sort of content.

Q345 **Chair:** I am pleased that Facebook has now agreed to conduct an analysis of whether Russian agencies were involved in distributing content on the platform that was linked to the Brexit referendum in the UK and other elections too. Just to note, however, as we are in the hearing now, that we expect that that will be based not just on whether rouble payments were made for promotional advertising but on an analysis of the origin of the content and whether it is likely to have come from a Russian agency due to the nature of that content and the way it was distributed. Rouble payments for advertising are one tool that could be used, but there are many, and we hope that the analysis will be wider than just that.

Simon Milner: Perhaps I could add to that. I can tell you that, as I explained previously in my letter to you, a second investigation is under way. It is not yet completed, but we expect to be able to report back to the Committee by the end of February on the results of that investigation. We will be prepared to share with you—possibly in private, because we do not want to tip off the bad actors about how we do those investigations—exactly how that work was undertaken.

Q346 **Chair:** Thank you. We look forward to seeing that at the end of the month. Monika Bickert, to what extent do you feel that the company has a responsibility to ensure that your customers are protected by knowing the source of the information that they see—where it is coming from and who is producing it?

Monika Bickert: We feel very responsible for letting our community know, first that they are in a safe community—some of the questions that were put to the previous panel spoke to questions of safety, and we can speak a bit about that. That is a huge priority. It is also a big priority for us to help people to connect with authentic information. From talking to our community, we know that that is something they care about, and it is something we care about.

We are investing a lot, and not just in the policies for keeping our communities safe, which tend to be simple in the sense that it is fairly black and white—it either crosses the line into something that is unsafe and remove it, or it doesn't and we leave it up. In the area of fake news, which is very different, we have developed a four-pronged approach where we are trying to make sure that when people come to connect with news

on Facebook, it is reliable news and they have the ability to make informed decisions.

Q347 Chair: Do you feel that not just with news services but with other information that people might share—community pages, pages that may have political content—it should be really clear to people where those pages are being administered from? If a community page is giving me information about Kent in England, where I live, it should be clear that that is being run by someone who lives in England, and not by someone who lives in another country.

Monika Bickert: There is a spectrum of different types of what people might call news or information. One of the things that has been notable about social media is that it has given a voice to many people who want to share what is happening to them, especially in regions of the world where traditional news media outlets do not necessarily reach, or do not reach frequently. There is a spectrum of information. Our job is to try to make sure that people can connect with reliable information and make their own decisions about what they want to see.

Q348 Chair: Isn't one of the ways in which you empower people to make those decisions for them to understand where the content is being created and whether the person creating that content is actually who they are pretending to be?

Monika Bickert: There are a couple of things that we try to increase transparency of source information. One thing that we do on Facebook, which is distinct from many other services, is have a policy that requires that you use your real name. If you think about the worst types of false news—the financially motivated spam that has links and that takes people to ad farms—that sort of content is very typically propagated by fake accounts. That transparency requirement in our policies is very important for removing those accounts.

Another thing we do is look at using context as a way to inform people about their news source. This is something that we are testing right now. When people see information from a news source, if there is any signal to us that that news source might be unreliable, they can click on a little icon recalling the article's context. We released this in November 2017. From that icon they are taken to information that is taken from across the internet about this source and its reliability.

Q349 Chair: On the point about labelling and transparency, people do set up fake accounts. You have identified fake accounts that have been taken down and there are very many of them.

Monika Bickert: We remove them every day.

Q350 Chair: So it is possible to do. But what would be the harm in making really clear the origin, at least, of material as you see it and consume it on the platform? Seeing where it is being created from may be a very big signal to you about whether it is a source of information that you should trust.

Monika Bickert: Yes, and because of our real name policy, people can see who is publishing that information.

Q351 **Chair:** As long as they are using their real name.

Monika Bickert: If they are using their real name. If they are not, we take that account down.

Q352 **Chair:** But people may not know and it could be there for a while. People do not necessarily know the country from where the page is being administered.

Monika Bickert: You are very right that we do not catch every fake account at its inception. We find and remove many fake accounts every day. This is also an area of tremendous technical investment for us and we have got a lot better at it. In the run up to the French, German and UK elections, we used our technical tools to remove thousands of fake accounts—not that those were necessarily related to spreading disinformation or information about the election. We are using those technical tools to reduce the chance that they might be used to spread disinformation.

Q353 **Chair:** Finally, I want to ask you the same question that I asked YouTube earlier on. What sort of proportion of your revenues from Facebook do you reinvest in identifying this kind of bad content?

Monika Bickert: I can't give you a revenue percentage, but what I can say is that we just put out an earnings call where our CEO said that more than 14,000 people at Facebook are working on safety and security issues. That includes the engineers who are working on the technical systems to identify fake accounts, terror propaganda or other violating materials. It also includes the work of our content reviewers who are looking at the sort of content that has been reported to us and removing it if it crosses the line.

Q354 **Chair:** YouTube said that they were spending tens of millions of dollars on this sort of work. How much is Facebook's investment in money terms?

Monika Bickert: I wouldn't have a number to give you. This is something that is such a priority around the company that more than 14,000 people are working on it. It is not just a question of sponsoring a certain programme. This is these people's jobs.

Q355 **Chair:** Do you know what the ad revenue is for Facebook for a year?

Monika Bickert: I believe our revenue for the last quarter was around \$13 billion.

Q356 **Paul Farrelly:** All my questions are for you, Monika Bickert, because I do not easily get the opportunity to interview you in London. Therefore, hopefully the UK taxpayer will get the best value from my travel. You have just mentioned thousands of fake accounts in connection with the US election. All I have seen so far are the 470 connected with advertising, so could you elaborate on that statement?



HOUSE OF COMMONS

Monika Bickert: My statement was that, in the run-up to the French, German and UK elections, we removed thousands of accounts using these enhanced fake accounts technical tools. We have been investing in this area for a long time. I want to be clear that the real name policy is not a new policy and using technical tools to find fake accounts is not new. I have been at the company for years and we have been doing this for years, but we have had a significant advancement in the past year. That has allowed us to remove those thousands of accounts in the run-up to those three European elections.

Q357 **Paul Farrelly:** So the thousands were not in relation to the US election?

Monika Bickert: That is right.

Q358 **Paul Farrelly:** So why only 470 with the US election? Who is better at French—you or the Russians?—than English?

Monika Bickert: I can refer you to the comments that were put to the Committee by my colleague, Colin Stretch. That is really the best place to find the information about that.

Q359 **Paul Farrelly:** You have not briefed yourself before this session?

Monika Bickert: That is part of an ongoing investigation. We are certainly co-operating with the relevant authorities there. We have given information and that is the best source of information for that.

Q360 **Paul Farrelly:** Do you do these sweeps in relation to specific events such as elections, or do you do them all the time?

Monika Bickert: We are doing them all the time.

Q361 **Paul Farrelly:** How many thousands or hundreds of thousands of fake accounts have you suspended that have no connection with specific events such as elections?

Monika Bickert: Without knowing precisely, that is probably the most common scenario. We remove many false accounts every day. Many of those are created, for instance, for the purpose of sending out spam links or engaging in other bad behaviour. We catch some of them at the time of creation and we stop them from creating the account. Others we can remove quickly, after identifying them using certain signals that our technical tools recognise.

Q362 **Paul Farrelly:** If you cannot give us the numbers now, could you provide us with a briefing as a follow-up?

Monika Bickert: We are certainly happy to follow up, but just to set expectations, when you remove accounts very quickly from creation, you do not necessarily know what the purpose of those accounts might have been. Any accounts that we find have been created en masse, or created where there are signals that they are not being accurate in their name or engaging in a false way, we remove regardless of why they came to Facebook.



HOUSE OF COMMONS

Q363 **Paul Farrelly:** We will ask you for some statistics afterwards. In relation to the 470 and the \$100,000, or a roubles equivalent, why did you accept that money in the first place?

Monika Bickert: With regard to the US inquiry, that is part of an ongoing investigation. We are continuing to co-operate with US authorities and that investigation will continue. For now I can simply refer you to Mr Stretch's comments to the Senate Judiciary, and to the Senate Intelligence and House Intelligence Committees.

Q364 **Paul Farrelly:** The question is unanswered. Why did you accept the money in the first place? Do you not make any efforts to know your user or your advertiser?

Monika Bickert: We do. With regard to our systems generally, we can speak about ads and also user-generated content. When it comes to ads, every ad that comes to Facebook is reviewed either by automated or manual review before it goes live. An important component of advertising on social media is that it happens quickly, so we try to use these systems to find things such as bad content—an ad that might have a certain word in it—that would suggest that we should take the time to review it before it goes live. There are a combination of signals that might lead us to take that sort of manual review. After the ads go live, the review does not stop. We look at signals like—

Q365 **Paul Farrelly:** Hold on. You have taken money in the first place. If you do not know your advertiser or your user, how can you be sure that you are not in breach of international sanctions or money laundering regulations? What responsibility do you take?

Monika Bickert: We have a team that works very hard to make sure that, when it comes to taking money, we are complying with all laws such as those around sanctioned individuals and countries. As I mentioned briefly before, we also have a policy that requires accounts to be authentic, and the advertisers must have an account before they can purchase an ad on Facebook. Our advertising is a self-service model. What that means is that if you use Facebook and you have an account, you can run an advertisement. If you do so, that ad will be reviewed in some fashion before it goes live, and then we will look at additional signals after the ad goes live, including how people are interacting with that ad. Are they x-ing it out, are they reporting to us or are there other signals that might suggest it warrants further review?

Q366 **Paul Farrelly:** My son, who is now 19, opened his Facebook account at the age of nine, and has since been unable to change his birth date. Your checks on who signs up stumble at the first block. Regarding the integrity of Facebook, it is important for users and advertisers to be reasonably confident that they are dealing with real people, not fake people or bots. What can you do to improve your game in making sure that people are dealing with real people and not fakes?

Monika Bickert: That is a very, very important part of what we do. When people come to Facebook, they expect that they are interacting with real



HOUSE OF COMMONS

people. That is fundamental; it is a cornerstone of our service and it is something we invest very deeply in. Our minimum age for people to come to Facebook is 13, and we have automated systems to try to detect when a person is putting in a false birth date. As Simon said, some things we can follow up in private because we do not want people to game those systems. We have systems that restrict people from changing their birth dates, for a similar reason. It does not always function perfectly, but we catch many people that way who might try to come online earlier in violation of our policy.

Detecting the fake accounts is something we are constantly investing in and constantly improving. Where we have safety-related policies such as not allowing people to share extremist content or not allowing people to bully others, some of that can be done by technical tools—like Google and YouTube, we are investing in that and seeing real gains from it—but a lot of it is also contextual and requires human review. That is also true with fake accounts. There are technical tools that allow us to identify that an account may be fake, and sometimes it is very simple to say it is, and we remove it. Other times it is not as simple, and we will send it to one of our reviewers, who will look at it and say, “Actually, it appears that this is real,” or, “We’re going to ask this person for identification,” or, “It does appear that this is fake,” and they will remove it.

Q367 Paul Farrelly: Final question: if I come to Facebook and complain that my identity has been stolen, how long in terms of your policy would it take for that page, post, site—whatever you call it—to be removed by Facebook so that I can be satisfied that you are acting?

Monika Bickert: The vast majority of all complaints that we get from people about violating content or violating accounts and that sort of thing are reviewed within 24 hours. If there is something we think is safety-related, it goes to the front of the queue. If your account were hacked, for instance, and you had that sort of concern, that is something we would respond to very promptly. If you are seeing an impostor account, that is also something we would attempt to respond to very quickly. Impostor accounts, obviously, could really wreak havoc.

I will say that that is not always a simple inquiry. I have looked at our reviewers who do this work, and I have watched what they do. Often what it entails is that you have one account and another account, and the second person has said that the other is an impostor. You have to look not only at the dates the accounts were created and the content within the accounts to try to figure out which one is right, but you might have to ask for the upload of forms of identification to confirm. We try hard to resolve all of those within 24 hours, but we do not always hit that mark.

Paul Farrelly: Thank you. I think we all think you could do much, much better.

Q368 Julie Elliott: Before I get to what I am going to ask, I want to come back to what you said about taking a large number of fake accounts down. Have you done an analysis of which countries the fake accounts come

from in the main?

Monika Bickert: You mean when I said in advance of the French election and the German election?

Julie Elliott: Yes.

Monika Bickert: There are various signals that our team uses, and we can follow up privately and walk through some of those signals.

Julie Elliott: Right. Certainly in the UK, there are significant implications around electoral law for people who are trying to influence our elections from outside, because all materials in elections are very regulated. So where the fake accounts are from is very important, so I would appreciate that, if you could find out.

Monika Bickert: Absolutely.

Q369 **Julie Elliott:** In the UK we had a very controversial referendum some 18 months ago. Do you believe that disinformation campaigns using your platform played a role in that referendum?

Simon Milner: I am more expert on that issue. We understand why people are concerned about it. Concerns have been raised in Parliament, and it is one of the reasons why, when the Electoral Commission approached us around October last year and said, "We would like you to assess whether or not there was indeed misinformation coming from another country—particularly Russia—associated with the referendum," we were very keen to co-operate on that. We reported some findings from an initial investigation to the Chair of the Committee in December. I think it is fair to say he felt we had not done enough work and, having reflected on that, that is why we are now undertaking more work. So the answer is: we won't be able to tell you until that work is completed, but we are committed to telling the Committee the outcome of that at the end of February.

The one other thing I would say is, unlike the US election, we have still not been furnished with any intelligence reports from the UK authorities to suggest that there was direct Russian interference using Facebook in the Brexit referendum. That is quite different from the US, where there is an intelligence report demonstrating this.

Q370 **Julie Elliott:** If we put Russia to one side, obviously the referendum was about our membership of the European Union—or not—going forward. So countries within the EU other than the UK probably had more of a vested interest. Are you looking at any countries, or just at the Russian influence?

Simon Milner: We have been specifically asked to look at Russia. I have not been aware of any parliamentary debate, any newspaper stories or anything else suggesting that countries apart from Russia might have been doing this. If you have got something you want to furnish to us—

Q371 **Julie Elliott:** It is not whether we are suggesting anything. I am asking,

as a company, are you looking at that?

Simon Milner: No, we are not.

Q372 **Julie Elliott:** You are not. There was, very clearly, misleading information on Facebook during that referendum, because, anecdotally, every time you knocked on a door and somebody would tell you something—it would be absolute nonsense—you would say, “Where has that information come from?” and they would say, “Facebook”. That was every time. What as a company can you do to stop that proliferation of false information getting shared and re-shared? I think people were buying into it because those sharing it were people they knew and trusted, but goodness only knows what the source was.

Monika Bickert: Thank you for the question. I want to be clear that we do not accept that sort of proliferation of false content on Facebook. At the same time, I want to make sure that we are distinguishing between the sorts of content like extremist content and bullying content, where there is a bright line of leave it up/take it down, versus fake news.

When we hear the term “fake news”, as you point out, many people will share stories of things they have seen online, and those range everything from the financially motivated spammers linking off-site to ad farms, which is the most common type of false news we see on the platform, down the line to the sensationalist headlines where the underlying story may be based in truth and perhaps given a spin or using certain words to get people to click on a headline. We cannot have one policy that addresses all of that; we have to have a more nuanced approach.

Since this has become a topic of interest and concern, especially over the course of the past year and a half or so, we have developed a four-pronged approach, which I am happy to walk through. If we think about why it does not make sense to have just one approach, you can take this to the point where you say, “What if we had a policy on Facebook where we required people to only post what they knew to be true and accurate?” Of course, that would lead to a place where not only was it unenforceable, because we would not know whether something an individual posts about his or her observations and daily life was true or false; it would also really inhibit the type of speech that we all engage in on a daily basis, including predicting, speculating or sharing opinions on things. We have to tread carefully.

So the first thing we do is remove false accounts and known bad actors from that far end of the spectrum that I talked about. If we do it well—we are not perfect at it—that takes care of a lot of that content.

The next thing we try to do is disrupt the financial incentives for these sorts of actors to come to Facebook. I mentioned earlier that they might post links that appeared to be exciting articles, but were clickbait articles that would take you off-site to an ad farm. We are getting better at detecting those ad farms and at detecting things like publishers inserting clickbait into their headlines or something that looks like a video that plays, when in fact the play button is a

ruse to get people to click on it and be taken to an ad farm. Our systems are detecting that and removing it.

The next thing we are trying to do is prioritise the visibility of content that is trustworthy, and specifically that is deemed trustworthy by our community. That is something we are testing right now, but we are very interested in it. We are trying to reduce the visibility of content where we have a reason to suspect that it is unreliable. People can report to us when they believe news to be fake. We are using a system of fact-checkers, which is growing. If we have an indication that news is fake, we reduce that visibility up to 80% in people's news feed.

The final thing we are doing, which I really want to underscore because I think it is so important in the long term for all of us, is trying to improve the ability of the broader community—not just people using Facebook, but journalists, policy makers, educators and parents—to fight false news by recognising it, distinguishing among news sources and being able to make those responsible choices. We are doing that in the UK. We have been working with young people—we have digital ambassadors in schools, talking among other things about how to recognise false news. In the run-up to the election, we ran Full Fact's top 10 tips for spotting fake news. That was not just something we did on Facebook; we actually went out in traditional media and published these top 10 tips to help people make responsible choices. It is also something that we have been doing over the past year with the National Literacy Trust in the UK. We are trying to research and understand the best ways to help people to recognise false news and disrupt it.

Q373 Julie Elliott: As a company, do you think you have any responsibility for the misinformation that is on your platform? I could take you to half a dozen sites that have absolute nonsense on them about the area I live in, but they get re-shared and everything. That is influencing people's voting behaviour and therefore influencing elections. I would say that that is having a very negative effect on the democratic process, which is between people of political parties who stand on their platforms. Do you think you have any responsibility to try to sort this mess out?

Monika Bickert: It is very important to us. I mentioned the broader initiatives, but I want to emphasise that we are engaging with the community and trying to learn more about what they are seeing and the way that this news manifests itself. We are exploring and testing different options—I have mentioned a few of them. In addition to the article context icon that you can click on to find information about the publisher, we are also testing a way of giving people information about related articles. If you are on Facebook and you see a news article about a particular topic, underneath that you will see a spectrum of related articles on the same topic, so that you know whether the article you are seeing fits with mainstream sources or not. We are looking at ways of incorporating brand logos, so that if you find content on Facebook—if it has surfaced to you—you know exactly who it is coming from and whether or not you recognise or trust that news source.

This is not an area where we are done; it is definitely an area where we are investing in learning and understanding and in testing different options.

Julie Elliott: Thank you.

Q374 **Chair:** I would just like to return to a couple of points that Simon Milner raised. For the record, my complaint about the analysis that Facebook did is that all Facebook had done when looking at the Brexit referendum in the UK was to go back to the accounts that had already been identified as part of the US Senate investigation, and look to see whether those self-same accounts had been active in the UK during the Brexit referendum—nothing else.

Simon Milner: Yes. I did not mean to dismiss what you said. You were very clear about your assessment, and your assessment and the views of other parliamentarians were a significant factor in why we are now doing another investigation.

Q375 **Chair:** And the view of the British Government, too. You also mentioned that the work that had been done in America on identifying the accounts that have been identified was based on intelligence that Facebook was given.

Simon Milner: I said that there was intelligence about Russian interference in attempts to interfere in the US presidential election, and that was a factor in the work that was done. I did not say that their accounts were developed. But I also—

Q376 **Chair:** Sorry, what do you mean by intelligence?

Simon Milner: An intelligence report was produced by the US authorities about attempted Russian interference in the US election. A similar report has not been produced by UK intelligence about attempts at Russian interference in the Brexit vote.

Q377 **Chair:** When I met yesterday with Senator Warner and Senator Burr, they were slightly surprised by that presentation of events. As far as they were concerned, there was no presentation of intelligence by the American Government or by them to Facebook, seeking to identify accounts that had been problematic. The accounts that were identified were simply a response by Facebook to pressure they were receiving from the Senate, which led them to do the bare minimum, which was simply to look for rouble payments that the site had received. This was not based on a dossier of intelligence that had been received externally.

That is why, when we first discussed this, I felt that Facebook, in response to a request from Parliament or the Government, should conduct its own research and not be reliant on other people giving you intelligence. Ultimately, this is your system and your platform, and you know it better than anyone else.

Simon Milner: Absolutely. I was not suggesting that therefore, this had a bearing on our ability to look properly at our systems. That is exactly what

we are doing. I was just explaining the wider context of the UK situation compared to the US situation.

Chair: Yes, but you were insinuating that there was a lack of intelligence in the UK that had existed in America, and that the absence of intelligence support from the UK meant that that work had not already been done. But in America, it was not intelligence reports from the Government that led to that work being done; it was pressure from Congress, which led to what they would see, I think, as the company doing the bare minimum. I am pleased that at least now Facebook is prepared to initiate that research in the UK, but I felt that the exchange that we had earlier did not necessarily give the clearest view of what had actually happened.

Q378 **Simon Hart:** I have more of the same. You recently announced that there would be 1,000 extra people to vet political advertisements. What problem had you identified that led you to make that decision?

Monika Bickert: Thank you for the question. When we looked at the advertisements that we identified in the wake of the US election, one of the questions we asked ourselves was "Are we doing enough to identify when ads might be coming from bad actors?" There are a couple of things that we asked ourselves. One is, "Are our policies in the right place about political ads?" We have actually decided to tighten those policies overall, for anybody who might run a political ad, to make sure that we are not allowing ads that inadvertently, for instance, contain hate speech. The second thing that we asked ourselves was "Are we looking—"

Q379 **Simon Hart:** This was in response to some evidence, presumably.

Monika Bickert: This was in response to us overtaking a broad look, in the wake of the US inquiry, at our advertising systems and saying, "How do we prepare going forward for what might happen during elections?" So it was not necessarily because of specific ads that we saw in the US-related investigation.

The second thing that we asked ourselves was, "Is our review process holistic enough?" We do not want one reviewer looking at the content of the ad and another reviewer looking at how the ad is being targeted; we want to make sure we have one source of information for understanding who is behind an ad, how that ad is being run, what is in the face of it and who it is targeting. We have made some structural changes to our review, and part of that requires investing more in our reviewers. That is why we are adding those additional people.

Q380 **Simon Hart:** You mentioned in response to an earlier question that there was a lack of intelligence, or no intelligence had been passed on to you, about the possibility that UK elections and referendums might have been impacted by political advertising from sources that had mischief in mind or were from other countries. Presumably, there was therefore no intelligence that that was not happening. You just do not know.

Monika Bickert: I am sorry; would you mind reframing the question?

Q381 **Simon Hart:** Mr Milner's point was rather a shrug of the shoulders; he



HOUSE OF COMMONS

said that no evidence was passed to them that anything went wrong with the EU referendum. There was presumably no evidence either way, so do you actually know whether there was a problem or not?

Simon Milner: That is exactly what we are looking into. Just to be clear, when the Electoral Commission approached us and said they would like us to look into this, we specifically asked whether they had any examples—had any member of the public, any parliamentarian or anybody from any campaign highlighted to them a page or an ad that was seen during that election that they felt was fishy? Perhaps it did not feel like it was coming from one of the official campaigns or came from somewhere that they just did not recognise?

The Electoral Commission was very clear that it had had no complaints. Again, I am not suggesting that there is nothing, but until we have completed this investigation, we will not know. What we have not had is information that has enabled us to target a particular page or a particular phenomenon from another source. To Mr Collins's point, that does not mean we are not looking very thoroughly; we absolutely are.

Q382 **Simon Hart:** One of the pitches that your company makes to us, as parliamentarians in the UK, is the effectiveness and cost-effectiveness of targeted advertising, so it clearly works, because you are encouraging us to do it most of the time. However, it is quite subliminal. It enables me to target people over the age of 65 with an interest in fishing in my constituency, for example. That is quite a subliminal method of political advertising. Are you happy, in pursuing that, that you are within electoral guidelines, particularly in relation to spending limits, which are obviously quite rigid in the UK?

Simon Milner: I had better take this. We have a very good relationship with the Electoral Commission, by the way. This is not the first time we have met them. We have worked with them on encouraging people to register to vote, reminding people it is election day and so on. We have a good relationship with them, and they are of course the ones who are principally responsible for determining how much money people have spent and where that money came from. That is not something we can necessarily see.

However, we absolutely agree with you that there is an issue with the transparency of political advertising. Can you see what your opponent in your constituency is saying to voters, and can you respond to that, if the advertising takes place on Facebook? That is one of the reasons why we are now rolling out a system of transparency around political advertising, such that, in due course—at the next general election in the UK, for instance—you will be able to see every ad that is being run by both the main campaign pages and by all candidates. If you want to see what ads they are running on Facebook, you can see them.

We are going to introduce a radical new level of transparency that has never been seen before in elections. We did not have that last time, and



HOUSE OF COMMONS

we recognised that it would be very valuable for people to address exactly the issue you are concerned about.

Q383 Simon Hart: You will have seen this week that the PM said that online platforms “are clearly no longer just passive hosts of the opinions of others” and that she believes it is right that we look at the liability of social media companies for illegal content on their sites. That reflects comments made in other parts of the world as well. Do you read that as being that the age of unregulated social media is actually coming to an end?

Simon Milner: I think one reads a number of things into that. We certainly do not think of ourselves as unregulated social media. Indeed, that is why Monika and her team are responsible for an incredibly extensive set of regulations that apply to our services.

Q384 Simon Hart: They are self-imposed. They are not independent regulations, are they?

Simon Milner: They are self-imposed, but I can give you an example of how they fit within a broader structure.

Q385 Simon Hart: That is not my point. I do not want to know that. We are talking about regulation that is accountable, democratic and transparent.

Simon Milner: It is certainly transparent, because all our rules are public. It is accountable in that I am here often, and in front of other parliamentary Committees, as are my colleagues, to answer your questions on that. It is accountable in that we let people know, when they complain about content, what decision we have made and why, and if someone’s content is removed, we let them know why we have done that. There are actually multiple layers of accountability.

Yes, there is not a single body from any country, or globally, that has said, “Okay, these are the rules that are applied to Facebook” from outside. It is very hard to see how that would work, but many countries have applied laws and rules that we take account of, and we ensure that people do not break those rules on Facebook.

Q386 Simon Hart: That strikes me as a pre-emptive way of opposing anything that might resemble state regulation. Do you fear outside regulation? If legislation did come forward that sought to apply a light-touch regulation to your business, in the light of some pretty overwhelming evidence, would you co-operate fully with that, or would you be minded to resist?

Monika Bickert: We would certainly want to be part of the dialogue, because from time to time we do see legislation that results in some unintended consequences that are not good for anybody. We would also note that our incentives are generally aligned with those of Governments. I was in the US Government as a federal criminal prosecutor for more than a decade before I came to Facebook. When I came to Facebook, it was right in step with a lot of the things I worked on as a prosecutor. All of the criminal behaviour and other things that we try to find and remove from our service were very much aligned with the incentives of policy makers.



HOUSE OF COMMONS

Like Simon said, we do have a process for complying with Government regulations. If a Government tell us that something is illegal on our service, the first thing we do is see if it violates our community standards. If it is something like terror propaganda, it would violate our standards, and so we would remove that globally. If it does not—let us say the law is about something that does not quite violate Facebook standards—we would look at the legal process and we might restrict the content in that country out of respect for the local law. So we do have a system in effect.

I would note that sometimes regulations can take us to a place—you have probably seen some of the commentary about the NetzDG law in Germany—where there will be broader societal concerns about content that we are removing and whether that line is in the right place. So we would want to be a part of the dialogue.

Q387 Rebecca Pow: Ms Bickert, isn't the clever part of Facebook that you get people to sign away their data and their rights to their data, and you find out everything you can about them?

Monika Bickert: No. Not only do we not sell people's data, which we never have and we never—

Q388 Rebecca Pow: I am not saying that you sell it. You get them to give you—they do not even necessarily know that they are giving you data. You are harvesting data about people every single time they use Facebook.

Monika Bickert: No, I would not characterise it that way. We are very clear in our data use policy how we do use data. We also allow people to see any information that we have. If you go to Facebook, you can go to "download your information"—it is literally called that—and you will see everything that Facebook has. So we are very transparent about how this works.

Q389 Rebecca Pow: But you gather all this data, whether they are playing a game, watching a video or whatever, working out what they are doing and gathering a profile, so then you can target them back with information that they might want. Isn't that the way you work?

Monika Bickert: The way targeted advertising works—we do allow that; this is how our advertising model works—is that advertisers will say they want to target this particular audience, for instance people who have liked their page, and we will provide an audience. We do not provide people's data, but we will provide an audience of people who fit that targeting criteria. That is the way the model works.

Q390 Rebecca Pow: You are gathering a massive amount of data about the whole of society. Isn't this a massive surveillance operation?

Monika Bickert: No, this is a system where people can come and communicate with one another and advertisers can target people based on the interests that they have expressed and the way that they have engaged on the site.



HOUSE OF COMMONS

Q391 **Rebecca Pow:** There is an article in *The Times* today in the UK and the headline is, "How to win 2 billion friends and destroy civilisation...the biggest surveillance enterprise in...history". It is by a very well-known journalist. He suggests that, whether Facebook likes it or not, you have got yourselves into this position.

You hold so much data about people that you are now very powerful, and you are completely unregulated.

Monika Bickert: We are definitely regulated in many different ways, and we are careful to—

Q392 **Rebecca Pow:** By yourselves?

Monika Bickert: No, by laws. Simon can speak to the UK.

Simon Milner: In respect of data, we are fundamentally regulated by European data protection law. Every single person on Facebook in the UK is covered by that data protection law, and we are absolutely accountable to it. It is completely wrong to suggest that, in respect of how we handle people's data, we are unregulated. The opposite is true.

Q393 **Rebecca Pow:** The *Times* article focuses on Channel 4, which is the biggest news outlet that has been using Facebook to post videos. It gets massive hits—2 billion views. You have placed adverts around those videos, sensibly, because you know they are going to get loads and loads of hits. I believe Channel 4 asked whether you might share the data of who was looking at those adverts so that they could get some information about their own audience. Will you share that data with them?

Monika Bickert: That is actually not how our ads work. Unlike some other services, we do not place ads next to specific pieces of content. That is just not how it works. You target an audience—

Q394 **Rebecca Pow:** That is not what they told me, but anyway.

Monika Bickert: It does not work that way. The way it works is that advertisers create a target audience, and people who meet that audience will see an advertisement, for instance in their news feed. The ad is not tied to a specific piece of content.

Q395 **Chair:** Just to be clear on that, there is ad space around the content. That is the point Rebecca Pow was making. People are buying an audience rather than saying, "I want to buy Channel 4 content," but there is advertising adjacent to the Channel 4 content.

Monika Bickert: That is correct.

Simon Milner: Yes.

Chair: That is the point that Rebecca Pow was making.

Rebecca Pow: Yes, that is exactly the point I was making.



HOUSE OF COMMONS

Monika Bickert: We do not run ads on pages—that is not something we do—but if somebody is in the audience for an advertiser, from the criteria the advertiser submits, they might see an ad in their news feed.

Q396 **Rebecca Pow:** I believe that you actually asked Channel 4 whether they could put some of your adverts in the middle of their videos, but they refused.

Monika Bickert: We may be speaking about something different here, which is an audience network, where there are ads off Facebook.

Simon Milner: It could well be. Perhaps we could follow up afterwards with you, because we are not aware of the specifics of the case.

Q397 **Rebecca Pow:** That would be useful. Sometimes, these videos are taken by other users, posted on and altered. I wonder where you stand on copyright issues in relation to things like that.

Monika Bickert: We have policies against anyone infringing others' intellectual property rights, and we have systems in place to detect that. Certainly, upon notice, we have a notice and takedown procedure.

Q398 **Rebecca Pow:** There has been a rather unpleasant incident of child pornography getting into the Channel 4 network through the Facebook Messenger service. That has actually been spread far and wide, and thousands of people have seen it. How did that happen? How did that get through your systems?

Monika Bickert: The way we deal with any sort of child sexual abuse imagery is that, first, we have systems in place to automatically identify such content. Those systems are not perfect. Perhaps I can explain the way they work. They primarily work on matching known images. For instance, we work with the National Centre for Missing and Exploited Children and the International Centre for Missing and Exploited Children, and if we become aware of an image of child pornography, we reduce that to a hash—basically a digital fingerprint—and we can stop it being uploaded to Facebook in the future.

We do not necessarily have the technical means to stop new images, or images of child exploitation that we have not seen before, being uploaded. As soon as we become aware of such an image, whether it is from someone reporting it to us or from law enforcement sending it to us, we will take steps to stop its dissemination and we will report it immediately to the proper authorities so they can take action.

Q399 **Rebecca Pow:** It actually took you 36 hours to take this down. How did that happen?

Simon Milner: I can talk to that, because I spoke with the editor of Channel 4 News about this incident. We ought to use this opportunity to say to anyone who is watching the live stream: if you ever come across a piece of content like this, do not share it. Report it to the IWF—that is what they are there for—and they will ensure it gets taken off the internet. Unfortunately, a lot of people shared this to condemn it, and were



HOUSE OF COMMONS

encouraged to do so. Every person who did that broke the law in the UK and in many other countries by doing so. Once we were made aware of it by "Channel 4 News" and once they told us about which in-boxes it was in, we were able to move very quickly to remove it from the in-boxes—to hash it, as Monika explained—and to prevent thousands more shares of that horrible material on Facebook.

Q400 Rebecca Pow: Doesn't this indicate just how powerful you have become and what a Pandora's box—to use the term my colleague used earlier—you have opened up? And doesn't it demonstrate that perhaps it is time for regulation and rules that set out the roles and responsibilities, and all of these things, because all of this is affecting our society, not least our children?

Simon Milner: I would encourage you to ask the IWF about this, because what they will tell you is that they see tiny, tiny amounts of this material being shared on platforms like Facebook. There are much more significant concerns about other aspects of the internet. Despite our big size—2 billion people using Facebook, 40 million people using it in the UK—these kinds of incidents are incredibly unusual. To suggest that this is a Pandora's box, or something we are not in control of, is completely wrong.

Q401 Rebecca Pow: My final question is, how long do you hold people's data for—the data you collect about them? Is there a limit?

Simon Milner: Yes—as long as they want us to. Mr Farrelly mentioned that his son has been on Facebook for 10 years. We have held his data for as long as he wants us to. If he wants to go back and look at data from when he first joined, he can do. But as soon as people decide they do not want to be on Facebook anymore, then we will remove that content. It is entirely up to individuals about how long they want us to hold it for.

Q402 Rebecca Pow: So he was a child when he signed up. How long do you hold data on children and do you remove it once they become an adult?

Simon Milner: It depends whether they want us to. It is entirely up to them. Remember, we are holding data for them that is very precious. A lot of that data may be photos and videos and family moments—important moments in their lives—and we are custodians of that for them. It is up to them to decide when and if that data is removed from Facebook. We will not make the decision for them.

Q403 Rebecca Pow: You do sound rather god-like, saying you are "custodians" of that data.

Q404 Chair: Let us move on. But just to be clear, when you say that it is for as long as they want, by that you mean for as long as they have a Facebook account.

Simon Milner: Yes. Exactly.

Q405 Chair: So as long as they have a Facebook account, you hold that data.

Simon Milner: Yes. We look after it for them.



HOUSE OF COMMONS

Monika Bickert: Although, for instance, if you removed a specific post from your account, then we would delete that in accordance with the relevant data deletion laws.

Chair: Okay. Thank you.

Q406 **Ian C. Lucas:** Mr Milner, earlier you said something very interesting, which I greatly welcome, concerning the election law, because in 2015 and 2017, and in the 2016 referendum, Facebook advertising was extremely important. In those elections, do you agree that it was not possible to establish, as a candidate, where another candidate's purchase of Facebook's advertising was bought from?

Simon Milner: I do not understand what you mean. Can you explain further? Can you give me an example in your constituency?

Q407 **Ian C. Lucas:** In my constituency, my opponent can purchase advertising from Facebook in the campaign. It is unlawful for someone to pay for that advertising from outside the UK. Also, all of the information has to be recorded within my particular district. To date, and as we stand today, it is impossible for me as a candidate to check where that advertising was bought from. Do you agree with that?

Simon Milner: Is that also true for the pamphlets that he paid for? Actually, was your opponent a man or a woman?

Q408 **Ian C. Lucas:** No—they have imprints on them.

Simon Milner: Right. But you still do not know how they were paid for, so I am not sure that it is any different from other forms of advertising—

Q409 **Chair:** Can you just answer the question?

Simon Milner: I do not understand what the question is.

Q410 **Ian C. Lucas:** The question is: can you assure me that foreign donors do not pay for campaign advertisements purchased in Britain today?

Simon Milner: I cannot assure you of that, no.

Q411 **Ian C. Lucas:** Do you hold that information at Facebook?

Simon Milner: No—let me try and think about the scenario. So, if somebody is buying adverts to run a campaign in your constituency during the election, we can see the account that has paid for the ads. We will not know where the money has come from to go into that account—

Q412 **Ian C. Lucas:** Do you know whether the account that has paid for the ads is from outside the UK?

Simon Milner: We will have information that will enable us to know who is paying for those ads, yes.

Q413 **Ian C. Lucas:** You know that is illegal—if someone paid from an account outside the UK.

Simon Milner: Yes, I am aware of that.



HOUSE OF COMMONS

Q414 **Ian C. Lucas:** You are aware of that. Do you prevent that from happening?

Simon Milner: We do not at the moment, but my understanding is that that is a matter for the Electoral Commission to investigate.

Q415 **Ian C. Lucas:** No, it is a matter for you. It is a matter for you, because—

Simon Milner: Actually, Mr Lucas, isn't it a matter for the person paying for the ad? They have to ensure that they comply with the law.

Q416 **Ian C. Lucas:** No, it is a matter for you. You are not complying with the law either, because you are facilitating an illegal act.

Simon Milner: I have never heard that analysis before. If you have something you can share with us that demonstrates that, I would be interested to see it.

Q417 **Ian C. Lucas:** See, this is the problem, Mr Milner. You have everything. You have all the information. We have none of it, because you will not show it to us.

Simon Milner: As I have explained, Mr Lucas, we are moving forward with new forms of political advertising transparency, which will enable you to have that information.

Q418 **Ian C. Lucas:** So there is a problem.

Simon Milner: Yes, but it would be wrong to suggest that we have suddenly determined that there is a problem on Facebook. There is a wider issue of public debate, which the ICO is looking at, the Electoral Commission is interested in, and perhaps Parliament will look at, and that is whether British election law—particularly around the transparency of spending—needs to be modernised for a different era of political advertising. Rather than wait for that, we want to do what we can to provide greater transparency in that area, and not because we have been told of particular examples of campaigns in your constituency or others where somebody was using foreign money to pay for their ads.

Q419 **Ian C. Lucas:** I am pleased that you have recognised that there is a problem, but it is a matter for Facebook. I welcome the assurances that you have given, but how do we check that the rules that you are announcing today are being complied with?

Simon Milner: You will, in a future election, be able to look at the page of your opponent, see what ads they are running and ask questions of them if you have concerns about where the money is coming from. The Electoral Commission has the power to require that person to provide information. Indeed, if you look at their annual report, as I am sure you have done, they regularly fine people for illegal election spending. They have those powers. We will provide information that will help you, but if you have a concern about that, you have to go to them to get that information.

Q420 **Ian C. Lucas:** What about advertisements from third parties—people who are not candidates in campaigns? Do you think that I would have access



HOUSE OF COMMONS

to sufficient information to check that?

Simon Milner: We are very hopeful that with our new transparency around electoral advertising—political advertising—you will have much more information than you have ever had before on the nature of advertising that has been deployed during an election.

Q421 **Chair:** It is extraordinary: if Facebook were a bank, and somebody was laundering money through it, the response to that would not be, “Well, that is a matter for the person who is laundering the money and for the authorities to stop them doing it. It is nothing to do with us. We are just a mere platform through which the laundering took place.” That bank would be closed down and people would face prosecution.

What you are describing here is the same attitude—it is up to the Electoral Commission to identify the person. Even though you know when money is being paid or linked to accounts outside a country, you do not detect it. We hear a lot about the systems, but they are not picking that up at all. Many people would find that astonishing.

I also want to ask Monika Bickert about the change in policy on election ads. That is a consequence of the passing of the Honest Ads Act in America, which requires disclosure about who is placing political advertising, is that right?

Monika Bickert: No, it is not a consequence of it. We are involved in conversations around that Bill, but right now it is a Bill. As Simon said, we don’t want to wait to see what Governments want us to do here. We have been looking at ways that we can get more transparent on our own.

Q422 **Chair:** As you say, it is a Bill rather than an Act, but it might have been helpful for Mr Milner to put that context in place as well. It sounded like this was something that had been started entirely on Facebook’s initiative, rather than as a consequence of a very public policy debate in this country.

Monika Bickert: We have certainly been part of that discussion, but this was something that we undertook voluntarily. After the 2016 election, we took a very hard look at how our advertising system works. Part of that is what I described in response to Mr Hart’s question—making sure that we are doing a better job of reviewing advertisements—but part of it is looking at how we can be more transparent. As I say, we are rolling out those initiatives now.

Q423 **Chair:** But you are doing this now to try to head off statutory regulation, not because you think it is a good idea.

Monika Bickert: No, this was something that we undertook voluntarily after the election in the US. We looked at our systems and tried to identify where we can do better. I should say that that is not just something that we do in this space; we do that across all our policies, all the time. Any time that we see a mistake, something that got through our system that should not have, or a new type of behaviour that we had not previously



HOUSE OF COMMONS

had a policy in place for, we update our policies. We are constantly looking to see how we can improve.

Q424 **Chair:** With respect to Mr Lucas's question, that was not what Mr Milner said. He was quite clear that the system is not picking up people from outside one country seeking to place political ads in another. It is not a question of you saying, "Well, we know it's not going on." What you are saying is, "It could be, and we're under no real obligation to call it out when we see it."

Simon Milner: But Mr Collins, we have not seen in the last general election, during the Brexit vote or during the 2015 general election, investigative journalism, for instance, that has led to the suggestion that lots of campaigns are going on, funded by outsiders.

Q425 **Chair:** You haven't looked, have you? That is the thing. You have not looked.

Simon Milner: Mr Collins, there is no suggestion that this is going on.

Q426 **Brendan O'Hara:** Just to move on a bit, given the controversy that has surrounded this issue in the last couple of years, can I ask whether Facebook has ever done an analysis of how profitable fake news or the deliberate dissemination of disinformation has been to the company? If not in dollars, is it: "not at all", "very little", "very profitable", or "hugely profitable"?

Monika Bickert: We are looking at research, which is ongoing. I cannot speak to the financial aspect of it, but I will say that that would not drive our decision. We have shown a willingness to find and remove false accounts. We remove the ad farms. If something violated our policies, whether or not we were gaining money from it, we would take it down.

Q427 **Brendan O'Hara:** I understand that that was the answer you gave a while ago. I just wondered whether there had been any analysis of just how profitable it has been for Facebook.

Monika Bickert: I am sorry, I cannot answer. I know that we have looked to understand how false news is manifesting itself. I cannot say that we have looked specifically at how much money has been gained, because it would not be a relevant factor in our decision to remove the content.

Q428 **Brendan O'Hara:** You say that it would not be a relevant factor. Do you then expect us to believe that, if the dissemination of fake news and false information had been a financial drain on the company, your reaction and your policies would have been exactly the same as they are?

Monika Bickert: We do think that, in the long term, because people come to Facebook for a safe environment where they can connect with reliable information, it is against our financial interests to have that sort of content on our site.

Q429 **Brendan O'Hara:** Do you think it has so far been against your financial interest to have that content on your site?



HOUSE OF COMMONS

Monika Bickert: We think that it is bad for our community, and in the long term people do not want to be in a place that they think is not safe, or where they think they cannot connect with reliable information. It goes again to the trust that our community has in us when they come to Facebook. That is critical to our business.

Q430 **Brendan O'Hara:** In the absence of any figure or an answer, I would guess that the propagation of fake news has been hugely profitable to Facebook. The more sensational the story, the more people are driven towards it and the more advertising surrounds it. Where does Facebook draw the line between the pursuit of profit and its social responsibility? If you have not crossed that line already, how far away do you think you are from crossing that line, and how will you know, by your own measure, when you have crossed that line?

Monika Bickert: I am very happy to speak to that, because my job at Facebook is to manage our policies. Money does not enter into it. We draw a line for every bad behaviour. As I explained earlier, fake news is different because there is a spectrum of behaviour, so there are many different types of lines we have to draw. We draw the line in our policies. If somebody crosses that line, even if it is a Facebook account that has many thousands of followers and is very popular, we remove it. If it is an advertisement that is set to run for a long duration and pay Facebook a lot of money, it does not matter. We remove it. If people cross that line, their content comes down.

Q431 **Brendan O'Hara:** You make it sound as if Facebook is an entirely benign organisation. Facebook is a massive—if not the biggest—player in the dissemination of information around the world. I suppose it goes back to what I was asking the folk from Google and YouTube about your role in helping to regulate that space. Where do you see that role? Are you happy with the regulatory framework? I am not talking about self-regulation; I am talking about a regulatory framework. Are you happy with the regulatory framework as it currently exists? Given that the debate is now happening, where do you see that debate going to the benefit not just of Facebook, but of society in general?

Monika Bickert: Thank you, Mr O'Hara. Because of the breadth of the sorts of content that is often called fake news—I want to be very careful—there are things, like my colleague from YouTube said earlier, that we would not necessarily call news at all. For the reasons I gave earlier, there are things where we should not be in the business of deciding whether they are fake or true. Our community would not want us, a private company, to be the arbiter of truth. There are, however, other types of content that we know are fake accounts intentionally trying to spread disinformation—those bad actors who are trying to send people to ad farms and that sort of content. Our incentives are very much aligned. We want to remove that content.

In terms of your question about where we are with regulations, we want to be part of that dialogue, especially because one thing I see when we talk to policy makers about these sorts of issues is that sometimes a



HOUSE OF COMMONS

regulatory solution might seem great in theory, but when we get together and talk about the practical ramifications, we and the policy makers will say, "Gosh, there are some unintended consequences here. We have to be very careful." An example of where that has happened recently is in Australia. There was a Government committee that recently considered this very issue. It was a long-standing inquiry, and they recently released their report. They said that they are not going to regulate for now and that they think collaboration is important. We can follow up with that.

Q432 **Brendan O'Hara:** But previously you have been accused of actively working against a legislative framework, particularly around electoral law.

Monika Bickert: I am sorry, in relation to what?

Q433 **Brendan O'Hara:** Electoral law. Why do you stand accused of that? Does that not make you directly complicit in what we are now experiencing?

Monika Bickert: We have reached out to electoral commissions since the US election in 2016. We came out and publicly said, "Here are some of the things that we are going to do to respond to the threat of false news." One of those is to actively engage with electoral commissions. That is an important part of getting this right.

Q434 **Jo Stevens:** Ms Bickert, my colleague just said that you are in fact the largest disseminator—some would argue publisher—of news in the world. You make de facto publishing decisions every day. We have heard you describe that. You design the algorithms and the algorithms then decide what we read and see on Facebook every day. Those algorithms are systematic. They are not objective and they inherit the biases of the people who are developing them. How many developers do you employ to design these algorithms?

Monika Bickert: I don't know the number of developers. I can tell you that we do have a team that works on those algorithms, and the algorithms are focused on surfacing content to people based on a variety of factors, including the recentness of a piece of content, how new it is, what the person—let us say you are using Facebook—tends to interact with and their friends and family. We recently came out with an announcement that we are prioritising content from friends and family. I don't know if anyone saw this announcement, but we said, "We know this is going to cause people to spend less time on Facebook. That's okay. We want people to have more meaningful interactions." Those are the sorts of things that go into that newsfeed algorithm.

Q435 **Jo Stevens:** You are the head of global policy management for Facebook and you do not know how many developers Facebook employs or has working for it.

Monika Bickert: For developers, no. I can speak broadly—

Q436 **Jo Stevens:** Roughly? Is it hundreds, thousands, a dozen?



HOUSE OF COMMONS

Monika Bickert: We have many engineers in the company, and some of them work on the team that refines that algorithm, and that is something they are constantly updating.

Q437 **Jo Stevens:** Would you be able to tell us afterwards privately how many are working on algorithms?

Monika Bickert: We are always happy to follow up.

Q438 **Jo Stevens:** Thank you. Your algorithms enable hyper-personalised content that is really finely grained to be directed and targeted towards specific individuals. Your advertising helps to do that, but most individuals who use Facebook do not even realise that you are doing that. They do not realise that what comes up on their Facebook—

Monika Bickert: The newsfeed.

Jo Stevens: Yes. What comes up there is what you are targeting towards them. So there is a huge power imbalance there, because you are controlling it and the person who is receiving it has no control over it. That kind of reminds me—if you will forgive the analogy—of an abusive relationship where coercive control is going on. Somebody is deciding what you see, hear and read, what you have access to. Can you see the parallels with that? Can you see why I would be concerned about that?

Monika Bickert: People have a lot of control over their newsfeed, but we are concerned that they sometimes do not understand quite how much control they have over their newsfeed. We have tried to make that a little more visible for people.

If you are on Facebook, you can go to a section called Newsfeed Preferences—you can also type it into the Help section—and select. The reason we have a newsfeed algorithm is that if you come to Facebook and you have a bunch of different friends and a bunch of different pages that you interact with, rather than show you 200 posts from the person who posted the 200 most recent things, we try to prioritise based on the factors that I mentioned before for relevance. However, if you wanted to turn that off, you can do that. You can go into your newsfeed preferences and say that you just want to see the content in reverse chronological order and we will provide that. But we try to make the newsfeed algorithm something that is sourcing and providing the information you want to see.

Q439 **Jo Stevens:** How do people get to know that they can do that? For someone who doesn't know in the first place that you are controlling what's seen on their timeline—you say there are mechanisms for them to stop that happening, but they don't know about them—how do people find out about them?

Monika Bickert: Our hope is that when people come to Facebook, they will take some time to explore the site and they will see, for instance, that there are important privacy settings. You have the ability on Facebook to control whose feeds you post; you can friend or unfriend people, you follow or unfollow. All of these settings—the newsfeed preferences are included within that—are really designed to give people control over their

experience. I want people to know about that. I think it makes people's experience on Facebook much better if they do know about that, so we try to make it fairly visible.

Q440 Jo Stevens: Going back to the developers, do you know what percentage of the developers are women?

Monika Bickert: I don't know that. We have issued a diversity report in the past, and we can follow up with you on that.

Q441 Jo Stevens: In the diversity report that I have seen, it says that there is a heavy male bias among developers—I think on the technology side. Only 35% of Facebook's total employees are women; on the technical side, only 19% are women. I talked about inherent biases before, so if you have predominantly male developers developing algorithms, there will be an inherently male bias in those algorithms, which means that what is being fed across your platform is inherently male in terms of its bias. Does that worry you?

Monika Bickert: We are certainly concerned about any type of bias, whether gender bias, racial bias or other forms of bias that could affect the way that work is done at our company. That includes working on algorithms; it also includes things like making decisions on enforcing our policies.

There are a couple of levels at which we try to address this. One is at the source. We have acknowledged that we want our workforce to be diverse and to reflect the community in which we work. We have initiatives ongoing, right now, to try to develop talent in under-represented communities, to try to recruit from those communities and to get those people into Facebook and help them succeed.

Q442 Jo Stevens: When did those initiatives start?

Monika Bickert: At least a couple of years—we can follow up with some details.

Q443 Jo Stevens: And how are you doing against the targets you set yourself? Or did you set yourself any targets?

Monika Bickert: We are improving. We have work to do. We think that we, along with many other companies in this society, are confronting these issues. It's going to take some time, but we are very committed to it.

Q444 Jo Stevens: So did you set yourself targets?

Monika Bickert: We do have more information that we can provide.

Simon Milner: Yes, we can provide more information.

Monika Bickert: I am sorry to say, I am not the expert on this, but I want to say that this is important to the company. One thing that employees go through at Facebook is called—this is another way we are tackling the issue—our "Managing Unconscious Bias" training. That is



HOUSE OF COMMONS

about the recognition that everyone has these biases in their mind. We have training that is designed to help people recognise that and adjust for it.

Finally, when it comes to development of our algorithms, enforcement of our policies, we do have checks in place. For instance, with the way that we enforce our policies, we try to make sure that those policies are sufficiently granular that they don't leave room for somebody to interpret seeing one way or another based on these biases. And we do quality audits of each reviewer to make sure that we are getting that right.

All of this is to say, Ms Stevens, are we perfect on this? No. We have work to do. But it is something we care very deeply about.

Q445 Jo Stevens: My final question is on another issue. What I have heard this morning, from both you and the panel before, is that you recognise that there is a lot of work to be done; you talk about prioritisation to make your platform safer and better. Why not make data available across platforms to academics or to a commission, so that they could independently research and work with you to help to provide solutions?

Monika Bickert: First of all, we are working across the industry. This is something that we do fairly commonly in safety areas, and we are now doing it when we think about tackling fake news and elections integrity. So far as working with research organisations, that is something that we are doing—in fact, we are doing it in the UK with Media Trust. We will continue to find ways of partnering with these organisations.

Of course, we have to be careful. We have data privacy regulation that limits the data we can share, and we want to make sure that we always comply with that and that we are consistent with our terms and the laws. However, it is something that we do.

Q446 Christian Matheson: I'd like to ask you about Facebook's relationship with Cambridge Analytica. How would you characterise that relationship?

Simon Milner: Right now, we have some colleagues who are meeting with the Information Commissioner's Office to talk about that. As you may know, the ICO are undertaking an inquiry into the issue of political polling, and one question they have asked us is about our relationship with Cambridge Analytica. I am afraid I am not the expert on that, but I would be happy to follow up on it. We can provide to you the information that we provide to the ICO, who have asked us the same question.

Q447 Christian Matheson: Well, let's see how we get on with the questions and if you are able to answer them. Have you ever passed any user information over to Cambridge Analytica or any of its associated companies?

Simon Milner: No.

Q448 Christian Matheson: But they do hold a large chunk of Facebook's user data, don't they?



HOUSE OF COMMONS

Simon Milner: No. They may have lots of data, but it will not be Facebook user data. It may be data about people who are on Facebook that they have gathered themselves, but it is not data that we have provided.

Q449 **Christian Matheson:** How will they have gathered that data from users on Facebook?

Simon Milner: There can be all kinds of things that these organisations do. I think what data they have would be a good question to ask them, rather than us. We have no insight on that.

Q450 **Christian Matheson:** We may well do in the future and will see how that inquiry progresses. Is it the case that third-party users—app users, or whatever we might call it—can ask for a Facebook user's data and then pull that data off Facebook and bank it?

Simon Milner: Yes, that is part of the platform policies that we have. Perhaps Ms Bickert can talk to how that works.

Q451 **Christian Matheson:** And is it also the case that, when I, for example, agree to give my data, it also takes my friends' data as well?

Monika Bickert: No. We have policies—called our platform policies or our developer policies—that govern how these applications can use Facebook data. The way it works is essentially that they have to tell each Facebook user who is going to use their app what data they are requesting from them. An example might be their home town and their email address. They have to give a possibility to opt out of giving any non-necessary data. You see the elements of data that they require to run their app and which they don't, and you make those choices. You do not take data that is beyond you or personal data belonging to your friends with you.

Once you have the data that the user has agreed to give you, you have certain responsibilities under our platform policies. For instance, if you were to turn around and give that data to some third party or to sell that data or to engage in any of that sort of activity, that would violate our policies. If we found out about that, we would certainly enforce upon it.

Q452 **Christian Matheson:** And how might you find out about it?

Monika Bickert: We can follow up on this privately, but there are things that we do to try to discover that sort of behaviour. Of course, if somebody raised a flag to us, we would investigate.

Q453 **Christian Matheson:** And when was the last time that such an incident happened?

Monika Bickert: I don't have an answer for you on that; I am sorry.

Q454 **Christian Matheson:** How often does such an incident happen that you have to chase down on?

Monika Bickert: That's not common behaviour. We are very transparent with developers about our expectations if they are going to use our



HOUSE OF COMMONS

platform. We make clear that we expect them to comply, and we take steps to make sure that they do comply.

Q455 Christian Matheson: Has Facebook ever provided advisers to assist with political campaigns—the referendum campaign in the UK, for example—or embedded your advisers into those campaigns to advise on how they can optimise their micro-targeting?

Simon Milner: In the UK, we have two different teams. One team, which some of you may have interacted with, we call our politics and government team. That team sit within public policy and advise policy makers, parliamentarians and so on how to use our products. The team are focused on products that are free. They deal with how to set up a page, how to handle your inbox and how to manage your affairs. They also provide guidance now, and they did at the last general election, on how to keep safe on Facebook—how you deal with abusive behaviour on the platform, which is a very important concern for all of us. There is a separate team, who are involved in selling advertising. They obviously operate on a demand basis. If people want to buy advertising from them, they will help them to do that. Particularly during elections, when you tend to get an increase in political spend, those teams will work with those campaigns, but they are not embedded in campaigns.

Q456 Christian Matheson: During the referendum campaign, did you provide that kind of advertising service to either the remain side or the two principal leave campaigns?

Simon Milner: Yes, we did—to both.

Q457 Christian Matheson: And you did so with the Scottish independence referendum?

Simon Milner: Yes, to both.

Q458 Christian Matheson: You had a success story on your Government and Politics page recently. The only one recently was—I think you have replaced that story, which was a Scottish story, with one about a less controversial campaign in Finland. Is that—

Simon Milner: I am aware of this issue, and it's interesting that a news story has been written about it as if it is big news. Those kinds of case studies that we put out, we are often refreshing them, and really I think that this is a genuine case of somebody making a mountain out of a molehill.

Q459 Christian Matheson: So it's not because the whole thing has become a bit too hot all of a sudden.

Simon Milner: Absolutely not. We are very proud of the work that our teams do here to help campaigns that want to make use of our products to reach people with their campaigning messages. We think that is a fundamental part of how democracy works.

Q460 Christian Matheson: Let's just look at the ads that you have provided. If we look at the general election campaign of 2017, would you be able to



HOUSE OF COMMONS

provide all the adverts or would you be able to identify them and, if necessary, provide them? I am talking about all the adverts that were used to influence the course and nature of the campaign, including the dark ads, which are specifically targeted and only the sender and user can see.

Simon Milner: Mr Matheson, as I have explained, we are moving towards a system that will enable that. We are hoping that this part of it will be able to provide not just what ads are being shown right now, but an archive of political advertising. I'm afraid I am not able to tell you today how far back that will go. We think that what is particularly important is that we focus on moments of democracy that are happening now and which are to come. One cannot revisit previous moments of democracy and run them again; what we can do is focus our efforts on upcoming elections and referenda, and on ensuring we can help people to understand what is going on.

Q461 **Christian Matheson:** I think Mr Lucas asked this. You have identified that there has been a problem there.

Simon Milner: There is an understandable concern that campaigns and candidates have that they can't respond to or see the ads that their opponents are running. That is why we are introducing new levels of transparency to enable that.

Q462 **Christian Matheson:** Let's go back, then, to the referendum and to Cambridge Analytica. Do you hold information on how much money was spent on Facebook during the EU referendum campaign by Cambridge Analytica or its subsidiaries?

Simon Milner: I'm sure we will have some information in respect of the campaigns and about how much the different campaigns spent, but of course all those campaigns have provided that information to the Electoral Commission, and the information is accessible, I understand. I also understand, from our engagement with the Electoral Commission, that they will be producing a report quite soon on this issue.

Q463 **Christian Matheson:** Well, my understanding is that that is not quite the case, on the basis that with some of the funding, the allegation is that it was channelled through parties in Northern Ireland, where the rules on political reporting are different, and therefore that information is not forthcoming, but of course you might be able to provide it.

Simon Milner: As I say, we are co-operating with the Electoral Commission and with the ICO, who are looking separately at these issues. We are helping them as best we can. But I do think these are matters really for the Electoral Commission rather than for us.

Q464 **Christian Matheson:** Well, I will refer you back to the Chair's earlier statements on that.

Do you hold information regarding the content of online advertisements referring to the EU referendum, who they were sent to and who paid for them? Could you identify those?



HOUSE OF COMMONS

Simon Milner: That's one of the things we are doing in respect of this investigation that I referred to earlier. We are particularly focused on this issue of whether Russian-backed advertising was directed around the issues at stake during the referendum campaign. We will be reporting to the Committee in due course on that—before the end of the month.

Q465 **Christian Matheson:** Dr Martin Moore of Kings College London, looking at the automated bots campaign on Facebook, suggested that in the United States presidential election in 2016, there would have to be around 50,000 bot accounts. Would that be a figure that either of you would concur with? In order to send out the amount of postings that were—

Monika Bickert: No, I cannot confirm that number.

Q466 **Christian Matheson:** I will just ask one final question. Has Facebook ever been successfully hacked, Ms Bickert?

Monika Bickert: We have seen individual accounts by people that have been compromised. That is usually because somebody might give their password to somebody else.

Q467 **Christian Matheson:** No, but Facebook centrally and your back-end data store of user data—has that ever been compromised?

Monika Bickert: Not to my knowledge. I am not necessarily the person who would have that knowledge.

Q468 **Christian Matheson:** Would it be possible to hack in and alter—albeit for a matter of moments until your engineers discovered it—the underlying algorithms that are used to manage Facebook?

Monika Bickert: We have a dedicated security team that works to stop any unauthorised access to Facebook. They are identifying—the best way is to do that every day.

Simon Milner: Mr Matheson, we can send you a copy of a report published by our information security team around April last year on this very issue—something called “information operations,” which are attempts to disrupt and hack systems. We can send you that.

Paul Farrelly: I just have a couple of quick statistical questions. As of the end of last year, you had 2.2 billion monthly users. I think I was probably one of them, because I had to check up on someone who had said a few things about me. But I must say that I have cut down on Facebook, largely because of the Brexit issue, which is very topical. I was getting inundated by people at various points on the lunacy scale, from the abusive to the frankly deranged, which I can cope with, because I can just not read them and I have a fairly thick skin, but I cannot cope with friends of mine who read them and then say, “Do you realise what such and such person is saying about you?” At some point you have to stop. If I have to take that view, I just wonder, particularly with the addictive qualities, what effects Facebook's scale and success has on the mental health of children, quite frankly. Of the 2.2 billion users, just to inform your efforts,



HOUSE OF COMMONS

what is your best guesstimate within the organisation as to the percentage of those that are fake, non-genuine accounts?

Monika Bickert: Correct me if I am wrong, Simon, but I believe our estimates actually take that into account. In other words, we adjust for accounts that we think may be fake that we have not yet identified. We can certainly follow up on that. I do want to address—

Q469 **Paul Farrelly:** Let me just interrupt you. What is the scale of the adjustment then?

Monika Bickert: I do not know.

Q470 **Paul Farrelly:** Could you follow up?

Monika Bickert: Yes. I want to address what we do around harassment and hate speech. Rightfully so, it is a real area of concern for policy makers and others who are using Facebook. Here is our approach. For private individuals, we do not allow bullying of any sort. We would remove any content that we became aware of that was intentionally bullying somebody else. For public individuals—this includes elected officials—we would remove any sort of hate speech or direct threats. We do allow robust discussion and debate about public figures. Recognising exactly what you mentioned, that sometimes people will be saying things that are off-topic or irrelevant, one of the things we try to do is provide people with control, so that if it is on your page, you can control those sorts of comments and ensure that you can use your page the way you want to use it.

Q471 **Paul Farrelly:** I am aware of that; it is just simpler for me not to use it. A second statistical question: of the total number of posts—I haven't got a figure for that—or however you describe them on Facebook, from your internal efforts what percentage of those do you estimate are made by bot accounts?

Monika Bickert: I don't have an answer.

Simon Milner: I would expect almost none. That is not really an issue on Facebook because of the way Facebook works. It is more a matter for some other platforms you are about to hear from.

Q472 **Paul Farrelly:** Okay; and my final statistical question. I see that the UK ranks no. 10 in terms of the estimated number of users, and that India is actually ahead of the United States, with 250 million. What efforts do you make in your biggest user market to make sure that your platform is not used—in the world's biggest democracy, indeed—for electoral manipulation or to foment social unrest?

Monika Bickert: Just as we do with safety issues, we take a global approach. I spoke earlier about removing fake accounts. That is something that we are doing around the world. Outreach with industry, with Government commissions—that is something we do around the world. Something we have not talked about is what we are doing to prioritise good news and reliable journalism. Part of what we are doing is what we



HOUSE OF COMMONS

call the Facebook journalism project, which we launched last January, through which we have worked with more than 2,600 publishers to identify the ways that reliable news can best succeed on Facebook. This includes real product fixes, meaning for instance making it easier for news media organisations to attract subscribers or to have advertisements that work within their content. So this is something that we do—

Paul Farrelly: Perhaps we can follow up about your regional approaches in different countries.

Monika Bickert: Absolutely.

Q473 **Chair:** We have a session later on with some of the news media organisations so I am sure we will be able to ask them about that too. I just have a couple of clarification questions I want to ask before we finish. Going back to the discussion on Facebook developers, Mr Milner, you said that it was not true that developers had Facebook user data—but they had data about people on Facebook. I just wondered what the difference was between those two things.

Simon Milner: I suppose I was initially assuming that Mr Matheson was saying “Have you provided data? Has Facebook provided data to Cambridge Analytica or some other outside entity?” We do not provide your data to anyone without your permission; but the developer, the system that Ms Bickert was talking about, does allow people to decide, “I am prepared to let them have some of my Facebook data in order to get a service from them.” I am sorry if I implied—

Q474 **Chair:** I just wanted to be clear what you meant, because it was not clear at the time. A Facebook developer, then, has gathered data about Facebook users because they have interacted with the tools that they have created on the platform. If that Facebook user then decides to leave Facebook, does the developer keep the data that they have gathered?

Monika Bickert: No. That is in our policies. They actually have to delete the data once the person is no longer using the service. I would also note that you do not have to leave Facebook to make that decision. If you have interacted with an app on Facebook and you decide, “I don’t want to do that any more” you can go into your settings on Facebook and turn that off and reject that. Then they have to delete your data.

Q475 **Chair:** Monika Bickert, the issue of dark ads came up earlier—ads that are seen only by the person placing the advert, the person receiving it and Facebook. In the changes of policy you have put in place whereby you can see who the advertiser is and what other ads they have placed, would that include dark ads as well as display ads?

Monika Bickert: Yes. Any time you see an ad you will be able to see what is the page behind that ad, and what are the other ads that they are running.

Q476 **Chair:** Okay, but also if you want to see what ads a particular page has run, even if you are not in the audience for those ads, you will be able to see them.

Monika Bickert: Yes.

Q477 **Chair:** Finally, just on audience, this is a UK-specific story, but I will be interested in Monika Bickert's view on Facebook policy as a whole. *The Sunday Times* did an investigation in which it bought advertising space through Facebook for the 20 to 34-year-old audience—to advertise to them. They were told that the reach of that audience was 17 million people, but according to the UK national census there are only 12.3 million people in that age group in the country. So there was a disparity of 4.7 million people. To what extent is Facebook concerned that the audiences that it is selling do not actually tally with the people who are actually there?

Monika Bickert: I do not have an answer for you right now. I certainly want to make sure that we get to the bottom of that, and we will follow up with you on that.

Q478 **Chair:** It is quite a big discrepancy, a number like that. Of course many people in the advertising industry would say that if that discrepancy is real and continues it is basically fraud—mis-selling of an audience to an advertiser when a lot of that number could be fake accounts or people who have been wrongly ascribed to that category.

Monika Bickert: We certainly want to make sure that we are being honest and giving the right numbers to advertisers, so I will follow up with you on what happened there.

Chair: Okay, thank you very much. Monika Bickert and Simon Milner, thank you for your evidence.

Examination of witnesses

Witnesses: Carlos Monje, Director, Public Policy and Philanthropy, US and Canada, Twitter; and Nick Pickles, Head of Public Policy and Philanthropy, UK, Twitter.

Q479 **Chair:** Thank you very much for joining us for this final session. I am sorry that we are running slightly late, but hopefully we will not run any later than we already are. The members of the Committee will be returning, but I will start with some of my questions.

Mr Monje, you will be aware that the Committee has made repeated requests to Twitter for information about the activity of bot accounts and fake accounts linked to the Brexit referendum, particularly where those may be connected with Russian agencies. It is a similar request to that made by the United States Senate Select Committee on Intelligence. The only response we have had from Twitter is a re-analysis of some academic research that has already been done. It has not said whether it will, even in principle, conduct research into where accounts that have been politically active during British election campaigns are based. Are you able to give us an update on that? Will Facebook be able to supply the Committee with the information it has asked for?



HOUSE OF COMMONS

Carlos Monje: Thank you, Chairman, and thank you for having us today. On that question I would like to defer to Nick, my colleague from the UK, who has an update to give you.

Nick Pickles: Thank you, Carlos. Just to clarify—Twitter, not Facebook.

Q480 **Chair:** Did I say Facebook? Sorry.

Nick Pickles: Before my Facebook colleagues jump on me from behind. As Carlos said, and as I noted in my letters previously, we have been doing further investigation. I would like to read this, because it has various numbers in and I do not want to misread them.

Q481 **Chair:** Can I suggest, because of time—is it a very short statement?

Nick Pickles: It is two paragraphs.

Chair: Okay.

Nick Pickles: We can now update the Committee that our broader investigation, which we noted in previous letters, has identified a very small number of suspected internet research agency-linked accounts. During the referendum campaign, 49 such accounts were active, which represents less than 0.005% of the total number of accounts that tweeted about the referendum. Those accounts collectively posted 942 tweets, representing less than 0.02% of the total tweets posted about the referendum during the campaign. Those tweets cumulatively were retweeted 461 times and were liked 637 times. On average, that represents fewer than 10 likes per account and fewer than 13 retweets per account during the campaign, with most accounts receiving two or fewer likes and retweets. Those are very low levels of engagement.

Q482 **Chair:** What is the audience reach for those accounts?

Nick Pickles: As I say, we are looking at fewer than two retweets or likes per—

Q483 **Chair:** Yes, I heard what you said, but what is the audience for that? If you have a certain number of accounts being interacted with on a certain number of occasions and they themselves have people who follow them, what is the audience?

Nick Pickles: The engagement metrics we have been using in this investigation to understand how activities are being conducted on the platform have highlighted very low levels of engagement. Those engagement metrics will directly impact on the viewability of those accounts. If there are very low engagements, that would suggest there are very low views.

Q484 **Chair:** Fine, but what I would like you to tell us—although, as is always the case, some university academics may be able to work it out from the information you have just disclosed—is what the audience actually is. The number of accounts that are active is one piece of information, but the reach for that information is something else. What they are sharing will be of interest, too. Was it content that they had created, or was it links to



HOUSE OF COMMONS

other sources of information? That would be useful to know.

Thank you for giving us that update, but clearly we have other things that we want to follow up on and that we need more information about. We would be interested to know whether you have restricted your searches to certain already identified Internet Research Agency accounts, or whether you have done a trawl across the whole platform to look for accounts registered in Russia in particular that were active politically during the campaign.

We know that with Twitter there was evidence that large numbers of suspected bot accounts were created and then taken down after the referendum was over. That is why we are persistently asking for this information, but—with respect—I think we have got quite a prismic response to our requests.

Nick Pickles: I can touch briefly on that point. We were asked to look at City University research. One of the challenges we have is that these accounts were not identified by research by third parties. Twitter is an open platform. Our API can be used by universities and academics around the world, and it is, but unfortunately that does not give you the full picture. In some cases, people have identified accounts as suspected bots that have later been identified as real people. One of the things we do, which we work with academics very closely on, is ask people to bear in mind when they make assertions about the level of activity on Twitter that there may be cases in which those assertions are based on very active Twitter users who are real people and not bots. One of the dangers of using activity as a metric to identify bots is that you may misidentify prolific tweeters who are human. It is a benefit that researchers can use our platform, but it is a challenge for us that researchers cannot see the defensive mechanisms and the user data that we can.

Q485 Chair: Plenty of analysis has been done of the characteristics of suspicious bot activity. Twitter also knows where accounts are being operated from, and therefore you can easily detect the creation of accounts that have been operated in a different country that suddenly start tweeting about something happening in another location. That sort of activity is easy to spot on the site if you are looking for it. Carlos Monje, may I ask what co-operation was given to the US Senate investigation? Has the evidence of Russian-linked activity on Twitter just been extrapolated from the work that has been done on looking at Facebook pages, or is that separate intelligence that you have supplied to the Senate?

Carlos Monje: Thank you for that question. We are constantly monitoring our platform for any kind of activity that suggests that that is happening, from the Internet Research Agency in particular. We came across that information in a number of different ways. In 2015, there was a *New York Times* article about the activity on some of those accounts, and at that time we started actioning those—in June 2015.

In the course of the follow-up to the election, we got information from an external contractor named Qintel, which we share with other platforms.

That provided a seed of information that the IRA told us was related to this farm in St Petersburg. I believe that somewhere in the order of 100 accounts were turned over to us that we then looked at. Bit by bit, following more and more signals, we found accounts related to it. We have improved the information that we have provided to the public and—as you mentioned, Sir—to the US Senate and the House intel committee. Our number for the Internet Research Agency accounts that were active is now 3,814. We started suspending those accounts in 2015. All of them have been suspended. They are not functioning on our platform.

As you heard from our peer companies earlier today, we are very good at understanding what is on our platform, but sometimes it is important to have that partnership with third parties, with contractors, with civil society, with academics and with Government—law enforcement in particular—to help us figure out what we do not know and what we cannot see that is not on our platform. We are really good at tracking the connections on things that are on Twitter and sometimes we need some partnership on the rest of the picture.

Q486 Chair: Does Twitter believe that there are likely to be other farms, if you like, or agencies running fake accounts from countries like Russia? There has been a focus on one, but some say that if the level of activity people believe has been the case was carried out, that is probably too much to be done by one agency; there would probably be others as well.

Carlos Monje: I think we have to be humble in how we approach this challenge. To say that we have a full understanding of what is happening or what happened, we have continued to look and we are constantly looking for efforts to manipulate our platform. We are really good at stopping it, especially when it comes to the malicious automation side, but we will not stipulate that we will ever get to the full understanding.

Q487 Chair: Some of the evidence the Committee took when we started the oral evidence hearings in Westminster related to the referendum in Catalonia. Research done there suggested that there was not only Russian activity but activity from agencies based in Venezuela, messaging in Spanish. Has Twitter looked at that?

Carlos Monje: Nick Pickles, who is not only our UK lead but one of the leaders in the company when it comes to information quality, perhaps could address that more fulsomely than I can.

Nick Pickles: This is one of the challenges that Twitter presents, and an opportunity. Research is done and published. That particular research was not published in a journal. There is no underlying data and we have not received a formal communication of the findings. So sometimes it is very difficult for us to validate those external research findings, because all we have is the numbers at an aggregate level.

I would say one thing in response to your previous point, Chair, and the assertion that it is easy to identify very quickly where an account is operating on the internet and where someone is based. I was logging into my e-mail earlier on. As a standard corporate practice, we use a virtual

private network to communicate securely with our company. That took two clicks. As far as Google is concerned, I am not in DC right now, because my virtual private network is connecting somewhere else. So, the idea that companies have a simple view of how customers communicate with us—it may be routed through data centres, it may be routed through VPNs or it may be routed through users deliberately trying to hide their location. I want to caution the idea that somehow if it must be in one country, it absolutely is and there is certainty. All of this work at an information quality level is based on a variety of signals. We make probable decisions based on that, but they are very rarely absolutes.

Carlos Monje: Just to build on Nick's point, which I think is important, the geographic basis of where tweets are coming from and where users are is not always the strongest indication of what we use to action accounts based on violating our terms of service—which we take super-seriously. Even if Nick is dialling in from a VPN or a Tor browser or other way to obfuscate where he is coming from, if he breaks our rules—any of our rules—we will hold him accountable.

Q488 **Chair:** I understand the explanation you gave, Nick, that it is possible for people to hide where they are. But I imagine if we were talking to your advertising people, they would say, "It is quite easy to buy advertising on Twitter that targets people based on where they live." That would be one of the rudimentary requirements of a brand seeking to advertise on a platform like Twitter.

Nick Pickles: Only about 2% of our users share geolocation data in their tweets.

Chair: That is not what I asked, though.

Nick Pickles: That is often one of the indicators: people assume we have a geolocation. Now, someone may identify their country in their biography, and you may be able to infer it from the language they use. I am not saying that the research is not important; I am just saying that sometimes the conclusions reached do not match what we find as a company. We see that quite regularly: for example, some of the research on bots will identify people and some of the research will identify other manipulations on the platform that we were able to detect and prevent, but that is not visible—

Q489 **Chair:** So, if an advertiser came to Twitter and said, "I want to pay for promoting my messaging on Twitter. I want to target people in the state of Virginia", would you say, "We can't do that, because that is not the way we are set up", or, "Yes, we can"?

Carlos Monje: That is an excellent question—thank you. We work with our advertising clients to try to give them the best value for their money. We do not have as much information about our users as some of our peer companies. What we try to do is figure out the analogues that are trying to reach the audience they are trying to reach—followers of CNN, Fox News or the BBC? We do have a degree of geolocation within a country, or within a media market within a country, or within a media market within a



HOUSE OF COMMONS

country, but we do not over-exaggerate our precision on that. We do provide extremely good value to our advertisers, but again we are limited by some of the factors that Nick expressed.

Q490 **Chair:** But you would sell advertising based on location, even with those caveats?

Carlos Monje: It is one of the approaches, but often we find that follow graphs—people who are interested in certain subjects or search for certain issues—can sometimes be a better—

Q491 **Chair:** I understand that. Brands do not buy advertising just on location, but you could sell to an audience based on location?

Carlos Monje: We work with our advertisers to try to get them as close as they can to what they are aiming—

Q492 **Chair:** But yes?

Carlos Monje: By country, we do really well, like Nick said. About 2% of our users are actually geolocated at any moment, and so at any given time there is some uncertainty about where people are using it.

Q493 **Chair:** I want to be really clear. The question I am asking is, when selling advertising to someone who wanted to advertise on Twitter, would you sell an audience based in a location—a country?

Carlos Monje: Yes.

Q494 **Julian Knight:** You mentioned the figure of 3,814 fake accounts before—it has obviously gone up. We have had security briefings from people who have studied this subject since 2014, and they suggest that the number of fake accounts could be in the tens of thousands. In addition, the phenomenon that in many respects is probably more damaging is the means that are used to amplify these fake accounts and the disinformation. Often these accounts have a certain bona fide texture to them, but if you drill right down they have all the signatures of falsehood in that respect. Obviously you do not have the monopolistic position of Google, or the money of Facebook, but you seem to have an infestation of these types of accounts. Is this too much for you? Is this a little bit like the wild west?

Nick Pickles: I would prefer Mr Monje to answer this question.

Carlos Monje: I don't believe so. We are a smaller company. We have 3,700 or 3,800 employees worldwide. Google has more employees in their Dublin office than we have in our entire workforce. But we have been leaders on many fronts in utilising machine-learning AI to address some of the thorniest issues, and precisely the ones that you are talking about. I will give you one example, if I may: terrorism. Since the middle of 2015 we have taken 1.1 million accounts—foreign terrorist accounts—down. We identified, in the last six months, 93% of those before anyone else brought them to us—we used our technology to do that—and 75% of those before they tweeted once. We have incredible engineers and an incredible ability to address precisely the issue you are talking about: malicious automation.

We are now challenging 6.4 million accounts per week, and saying, “You’re acting weird. Can you verify that you’re a human being?” We measure our progress in weeks. That is a 60% improvement from October.

Q495 Julian Knight: But Mr Monje, what about the amplification of those accounts?

Carlos Monje: That is precisely what we are talking about: malicious automation, where you see a lot of people acting in a co-ordinated way to try to push something artificially. We have got really good at stopping that particular effort to manipulate our platform. We have protected our trends from that kind of automated interference since 2014.

One of the challenges that we see, and that Nick referred to, is that many of the folks who are investigating our platform are doing so through an API that we provide freely—we think it is in the public interest to do that. But a lot of things that we do on a day-to-day basis to action those accounts—to take down that malicious automation—is invisible to the API. The way we challenge accounts puts folks in a time-out, where they cannot be tweeting, but they are still visible to the API. Things like safe search, which is the standard default setting—

Q496 Julian Knight: One of the areas mentioned to us by policy makers over here is that you are exploring whether to allow users to flag tweets that contain false or misleading information. Where are you with that? How soon will we see that?

Nick Pickles: One more time?

Q497 Julian Knight: I had a document given to me about public actions that social media companies are undertaking. It mentions flagging tweets that contain false or misleading information. Where are we with that?

Nick Pickles: Sorry. Can you clarify?

Q498 Julian Knight: Are the users flagging tweets from other people, saying, “We believe this to be false or misleading”?

Nick Pickles: Sorry, I am curious as to the source of that because I am not aware of that being discussed.

Q499 Julian Knight: So that is not something you would actually consider.

Nick Pickles: It is more broad than that, because it goes to, first, some of the wider questions that the Committee is asking about what we would do with that information and the response required and, secondly, the likelihood of users gaming it to try and abusively remove people they disagree with.

Q500 Julian Knight: They are gaming you already.

Nick Pickles: This is one of the things we are very conscious of, as Carlos touched on.

Q501 Julian Knight: You have clearly not got the manpower, or person power, with 3,700 people. That is clearly the case. I appreciate that. Why don’t



HOUSE OF COMMONS

you allow more of your community to flag up these tweets that contain misinformation?

Nick Pickles: I want to clarify this. We are currently removing 6.4 million accounts every week for breaking our rules specifically around automation. That is a huge number of accounts. Are you asking us to remove content that is untrue?

Q502 **Julian Knight:** No. I am asking you, where you could potentially explore this area—I understood you were exploring this area—to allow other users, when they see a piece of misinformation, much of which is to damage political processes and stability within our own country, to flag that up as a warning to other users. Is that not something you would consider?

Carlos Monje: Through the whole sweep of your pretty incredible hearing today, you are hearing from a lot of different voices who are trying to approach the issue in different ways. Where there is a piece of misinformation that is against the law, we can action that extremely quickly. I will give you an example. During the 2016 election, near the end, in the US, there were a series of tweets that were spreading the idea that you could text or tweet to vote. It is a standard effort to mislead people and it is voter suppression, which is as old as electoral politics but has moved into the 21st century. That is against the law in the US and we are able to take that down extremely quickly. The truth reached eight times more Twitter users than the initial falsehood.

When it comes to plainly false information, the conversation that happens on Twitter at any given moment is a wonderful, cacophonous democratic debate. We are very cognisant of the things that we can control and the things that we cannot. We are trying to address malicious automation and the kind of things that can take a bad idea and spread it quickly. We are trying to address a lot of the things that Monika mentioned in the last panel and elevate voices that are credible, and then give our users a more sensitive understanding of who's talking and a broader view of who is actually speaking.

Q503 **Julian Knight:** In terms of outside help, considering your relatively small workforce, what about the need to better understand your users' information? What about what was mentioned earlier by my colleague, Ms Stevens, about Facebook and academic oversight, and allowing them in to see the information? I don't mean one or two, but a much more open and transparent means by which it can be done to the max.

Nick Pickles: I think that is absolutely true. I was at a conference in California last week with academics discussing this very topic. We as a company are currently hiring for two roles at our headquarters. Their jobs will be specifically to ensure those conversations are happening. Fortunately, you have already seen a lot of the information that academics have produced based on Twitter data. Twitter's API is open. Researchers already access it every day, from the Oxford Internet Institute to Yale and Harvard. There is a huge amount of research happening right now on Twitter data. We are absolutely committed to deepening that relationship

and learning more, but it is important to recognise that Twitter data is arguably the biggest source of social media data for the academic community as is. Only last week we expanded the amount of information available, increasing the ability to search for any tweet ever posted, rather than for a fixed time period. We can share the details of those changes with the Committee.

Q504 Chair: Based on your statement at the beginning, it would seem that these academics are probably more effective at finding problematic content than Twitter is itself.

Nick Pickles: One of the questions is about misinformation. If we are looking at how misinformation spreads and what kind of content spreads, one of the important issues is that Twitter is 280 characters, so often the content you are referring to might be a link to a newspaper or a video. Some of the information may be off-platform. We can absolutely learn how those networks are working, because it also informs the solution. One of the big challenges—a lot of research has been done on this—is that some of the solutions that have been proposed to educate users have actually had negative effects. They have created false trust, where that wasn't true, or more hyper-partisan feedback loops, where people see their own content and are worried that it will be removed so they share it faster. There is quite a lot of research going on there, and we want to learn from it, not just inform our approach to the behaviour but to improve the quality of the debate.

Carlos Monje: Last year, Twitter off-boarded Russia Today and Sputnik as advertisers, and devoted the money they put on the platform globally to exactly the kind of research that Nick talks about. We are unique among the big platforms in the amount of information we are giving out freely to the world. We know we can do better, and this is part of the answer.

Q505 Julie Elliott: You talked about the 6.4 million accounts that you have taken down. How many accounts do you have at the moment?

Nick Pickles: 330 million.

Q506 Julie Elliott: So it's a really tiny number in comparison with how many accounts you have got.

Nick Pickles: Well, your colleagues on the Science and Technology Committee thought it was a rather large number, so it depends on your perspective.

Carlos Monje: Just to clarify, we challenged those 6.4 million. It is all a question of certainty. The more certain we are that you are violating our terms of service, the more aggressive our actions are. Sometimes people in a moment of pique or excitement do things that are unusual—they tweet an unusual number of times. We will send them a note to say, "You're acting a little strange. Can you verify you're a human being?" If you're really into Twitter, you will pass that test, but if you are a bot farm, it is too expensive a thing to do.



Q507 Julie Elliott: You've got these millions and millions of accounts. The Twitter handles that people have often bear no relation to who they are and give you no idea who is following you and who is tweeting you. What are you doing to try to identify who people actually are?

Carlos Monje: There are two sides to that. Verification is an important tool that we have used for a long time. We are taking a hard look at that and are trying to figure out ways we can give our users more context about who the credible voices are on Twitter. As part of our effort to get ready for the 2018 elections in the US, we are working with the major parties to verify a larger number of candidates as a hedge against impersonation. Impersonation is against our rules. You can't say, "I am an MP from Manchester," if you are not. So that is against our rules; however, we do honour and respect certain voices that have to speak pseudonymously or anonymously. It's an important part of democratic debate; satire accounts are an important part of that.

It is important to note that, in various places in the world, there are freedom fighters, there are Christians in China, there are people in the middle east who are combatting ISIS and trying to promote counter-narratives against radicalisation. If they didn't post anonymously, their lives would be in danger, so we try to respect that. It is a complicated challenge, but as I said, we are taking a hard look at how we give people more confidence about who is speaking while also giving voice to those who are putting their lives in danger to communicate.

Q508 Julie Elliott: One of the things we are struggling with in the UK at the moment is the very abusive tweets that are coming to people in public life—particularly to women in public life. We have had debates in Parliament on that issue in recent weeks and months. Every day, as a woman MP, abusive tweets come at me on one thing or another. If I speak on certain issues, you just know. You kind of think, "Let's wait for this cyber-attack," and it comes. You never see who these people are. There might be a random name with numbers. There is nothing in the descriptor that says who they are. If you report it to the police, the police can't track down who they are. And so it goes on. There is usually a huge level of, not just attack in what they are saying, but disinformation as well, and these things just keep getting re-tweeted, and spiral and spiral. What as an organisation are you doing to try to stop that kind of thing happening?

Nick Pickles: I would pick up on the Committee for Standards in Public Life report. I met with the Committee twice to discuss its work and I am a former candidate myself, so this is something I have a personal interest in. I have many friends who have been in the electoral process.

First, when I joined the company at Twitter, our safety approach was in a very different place. I joined in 2014, just after the abuse we saw directed to Caroline Criado Perez and Stella Creasy particularly. Since then, our approach has massively improved. We have developed a lot of technology. We have strengthened our policies. We have built the tools in the platform to make them much easier. When I joined, I think it took 10 clicks to



HOUSE OF COMMONS

report content. It is now down to three or four. So we have made it easier for users. We allow users to report other content, not just content directed at themselves.

To give you one idea of the impact we think we are having—I can share the detailed breakdown—we take action roughly now on 10 times as many accounts every day as we did last year. In one year, we have been able to massively increase the number of accounts. You may have also seen that we took the decision in December to expand our policy on violent extremist groups, which led to Britain First and accounts affiliated with that organisation being removed from our platform, because we felt again that that was more we could do.

Something we are also doing now, which is relatively new—we are in America so I will use the American phrase—is the penalty box. When someone crosses the line, how can we change their behaviour to stop them crossing the line again? One thing we have done there is locking their account and saying, “This specific tweet has broken our rules—you must delete it before you come back to the platform.” Often we will say, “You must also give us a phone number.” Then they can come back on Twitter. I am just double checking the number, but of accounts that go through that process, about 65% only go through it once. So we think we are changing their behaviour.

That is not to say by any means that this is a done issue. Safety is still the company’s top priority. Jack Dorsey, our CEO, spoke about it being our top priority. But we think we are making progress and we are encouraged by the results of the work that we are doing.

Q509 Simon Hart: On that point, I am very encouraged by the shift that you have been able to achieve between 2014 and 2017, but I want to bring a case to your attention because it demonstrates the scale of the problem. A former colleague, Byron Davies, defending a majority of 25 in Gower in west Wales, was subject to a five-week campaign based on an accusation of a criminal act with which neither he nor his wife had ever been involved. Despite five weeks of effort to persuade Facebook and Twitter to do something about it, he was told that nothing could be done. Are you saying on the record now that that couldn’t happen again whenever the next election comes along?

Nick Pickles: No. I am familiar with this case. It is important to say that the story that was being circulated was published in a British regulated publication that is a member of IPSO. So the idea that this was an unfounded allegation that only lived on social media—

Q510 Simon Hart: I am not interested in what came up in print media. I am asking you about that demonstrably, provably untrue statement, which could arguably have cost his seat and altered the outcome of the election. Are you offering us a guarantee that that sort of thing couldn’t happen again? I don’t care what print media are doing; they are not in front of us.



HOUSE OF COMMONS

Nick Pickles: No, and you have heard from a previous witness, we are not the arbiters of truth. We are not going to remove content based on the fact that it is untrue. The one strength that Twitter has is that it is a hive of journalists, of citizens, of activists, correcting the record and information. But if—you know, we could have this same conversation about the £350 million number—

Q511 **Simon Hart:** I must be brief, because I am not scheduled to ask this, but the fact is that it qualifies, in that to a great extent it is abuse and/or intimidation. And it is untrue. But you are saying, “Well, actually, we’re going to do nothing about that.”

Nick Pickles: Context is important. If an account was created for the sole purpose of abusing somebody, that would cross our sole purpose abuse. If it was using language based on a protected characteristic and it is hateful conduct, we would remove it for that. But the truthfulness of a piece of information is not a ground upon which—

Q512 **Simon Hart:** So the deliberate use of your platform to distribute absolutely false and defamatory information can continue in election time, according to what you have just said.

Nick Pickles: More than that, I don’t think technology companies should be deciding what is true and what is not true during elections, which is what you are asking us to do. I think that is a very important principle that we should recognise. During the Brexit referendum, we have heard that there were similar claims made on both sides that people—

Q513 **Simon Hart:** So you are expecting a different level of regulation. If that had been a public broadcaster or a newspaper, legal action could have been taken, but you say that you should be outside those—that you should not be aligned in the same way as they are.

Nick Pickles: As I have said, the information originated in a currently regulated member of IPSO. The idea that there is some sort of distinction—we are not going to remove information based on truth, because I do not think that that is the role of technology companies. The idea that we are not regulated is not an accurate one, as you have heard from previous witnesses.

Q514 **Chair:** To be clear, for all the bad actors who are listening in, if you set up an account under an anonymous identity, you can disseminate as many lies as you like. That is not in breach of the community guidelines and Twitter will not act against it.

Nick Pickles: No, I did not say that.

Chair: I am just trying to extrapolate from what you have said to try to understand it.

Nick Pickles: I said very clearly that we take context into account. If you deliberately create an account to target an individual in a harassing and abusive way, we have rules in place—

Q515 **Chair:** What you are talking about there is harassment and abuse, which is clear, but we are talking about lies. Someone who decides to spread lies about someone else is not harassing them, intimidating them or inciting violence against them—they are just spreading lies. If they are using the anonymity of Twitter to do that, there is basically nothing that you will do about it as a company.

Carlos Monje: Anonymity on our platform is not a shield against breaking our terms of service.

Q516 **Chair:** But it isn't a breach in the terms of service, if it is just spreading lies.

Carlos Monje: It is conflating a number of things. Monika and Juniper mentioned earlier that fake news is an over-broad category. Labelling something as not true does not necessarily mean that the receiver of that information will discount it. We are focused on attacking the things that we can control: if something is illegal, it is against our terms of service. If it is telling people to tweet to vote, that is against our rules and we will take that down. Elevating credible voices is also incredibly important, and figuring out what are the voices—the BBC, *The Guardian*—that people—

Q517 **Chair:** Sorry, I want to finish this point before you move on, and I know other people want to come in as well. Telling lies on Twitter is not a breach of the community guidelines and would not require action to be taken against the account. That is what you are saying, isn't it?

Nick Pickles: If that's the only ground, and we didn't bring context into it, we do not have rules based on truth.

Q518 **Chair:** And obviously, people can be anonymous to protect their identity, so someone can use a false identity to spread lies about someone else. That could be reported to Twitter. They could be demonstrable lies—not a matter of opinion, but things that are demonstrably untrue—but that would not in and of itself require Twitter to take any action against them.

Nick Pickles: I don't want to sound like a broken record, but the context would be important.

Q519 **Chair:** I am not talking about the context. I understand what you said, but on this particular point: just spreading lies, not inciting violence or a campaign of harassment, is not a breach of the community guidelines. In that case, if I made a complaint to say, "Someone is spreading lies about me using a false identity. I can't take legal action against them because I don't know who they are," but it wasn't harassment, you would not do anything about it. Actually, what you said, Mr Pickles, is that you would feel that you were doing the right thing in not making a judgment on that case.

Nick Pickles: Yes, and I think that would be the same position across most big platforms. I do not think that is particularly unique.

Chair: In that case, that's probably why we are all here.



HOUSE OF COMMONS

Carlos Monje: I was looking this morning at the quote, “A lie can travel halfway across the world before the truth gets its shoes on.”

Chair: On Twitter.

Carlos Monje: It was attributed to Mark Twain—incorrectly. It was not actually him who said it. For all the stories of folks spreading misinformation on Twitter, there are as many stories of the truth coming out loud and clear. What we do as a company is try to figure what are the voices that resonate with people, that have credibility and that let the debate flow. Who would you put in that position to be the arbiter of the truth?

Chair: Throughout this inquiry on disinformation and fake news, it is quite clear that there are aspects of opinion which are up for debate—I appreciate that that is difficult—but sometimes there are downright lies. To use your misquoting of Mark Twain, from what you said, if you wanted to spread a lie halfway around the world before the truth got out, Twitter would be a pretty good medium to use.

Q520 **Ian C. Lucas:** Can I just be clear about the information that a Twitter account holder gives you when they open an account? Do you hold an address for that person?

Nick Pickles: Not a physical address, no.

Q521 **Ian C. Lucas:** So what kind of address do you hold?

Nick Pickles: We have an internet protocol address—an IP address.

Q522 **Ian C. Lucas:** So you don’t know what jurisdiction that person is in?

Nick Pickles: There will be a range of other information we might be given. A phone number might be one; an email address might be another. The user may share their location and their profile picture. There is a range of information that we may have.

Q523 **Ian C. Lucas:** So if I were disseminating information in a political campaign, I could use Twitter, regardless of which jurisdiction I was in in the world?

Q524 **Nick Pickles:** Yes.

Q525 **Ian C. Lucas:** And I could pay for advertisements on Twitter to assist targeting of that campaign?

Nick Pickles: Carlos is probably in the best position to talk about our advertising and transparency centre, which would cover that.

Carlos Monje: We shared with Congress the results of our 2016 election. It was nine advertisers that we believe—

Q526 **Ian C. Lucas:** Excuse me. I don’t want to be rude, Mr Monje. Can you just answer the question?



HOUSE OF COMMONS

Carlos Monje: About ads transparency? Yes. We were very happy last year to lead an industry-leading announcement when it comes to ads transparency and to develop a centre for electioneering ads, which are identified as ads that mention candidates. For all advertising, regardless of whether it mentions them, people will have to say what the account is behind that advertising, what the creative is—that's the term we use for the video or the advertising image—and an additional layer of information, including who is paying for the advertisement and demographic targeting. We didn't have to do this—

Q527 **Ian C. Lucas:** And where they are—where the source of the tweet is—and who they are?

Carlos Monje: Yes. Our centre is intended to enable advertisers to disclose where they are and precisely to be in compliance with election law.

Q528 **Ian C. Lucas:** In both the US and the UK, as I understand it, it is illegal to accept donations from outside the jurisdiction. Are you saying that your new rules will enable a candidate or the relevant authorities to be clear that a donation does not come from outside the jurisdiction?

Carlos Monje: It is our intention to develop a system so that the paying entity can disclose that information, and that the entire public can see that, and the historic spend as well.

Q529 **Ian C. Lucas:** So in none of the 2015 UK general election, the 2016 referendum, the 2017 general election and the 2016 presidential election was there a system in place to prevent someone from outside the jurisdiction illegally disseminating information in the election campaign?

Carlos Monje: That is correct. I will note that our investigation revealed a very small number: nine advertisers altogether. Two of those were the ones that we discussed: Russia Today and RT America, which is a branch of Russia Today. This is a news site that advertises on a lot of platforms. We are the only major platform to kick them off as advertisers.

Q530 **Ian C. Lucas:** But we come back to—

Carlos Monje: It wasn't a major—

Q531 **Ian C. Lucas:** Well, we don't know whether it was or wasn't. As I said to Facebook, you guys have all the information, we don't have any of it and you won't tell us what it is.

Nick Pickles: I want to clarify one point. You used the word "donation"—I think it was referenced in the previous hearing, relating to the exemption under Northern Irish law. If someone were to donate to a political campaign and that political campaign then spent that money on advertising, we wouldn't know the origin of the donation, which I think was your question previously. We would know who spent the money and who our customer was, but we wouldn't know the money had originally come from that source. I think that is rightly why the election—

Q532 **Ian C. Lucas:** It was actually my colleague who asked that particular



HOUSE OF COMMONS

question. I have some friends who are American politicians. If I illegally purchased advertising in the United States and disseminated it through Twitter, there would be no way of establishing that that was from the UK.

Carlos Monje: It is our intention, through the advertising centre, to address that. It is not just a question of building the technology to allow people to report, but to try to figure out ways that we can, at scale, actually make sure that people follow the rules. It is an important challenge. Again, I am very proud that Twitter was the first out the box on those reforms.

Q533 **Ian C. Lucas:** Could I just ask one brief question? Do you agree that Twitter is addictive?

Nick Pickles: One of the challenges of all technology and all tools and things that we use in our everyday lives is that different people react to them in different ways. Some people are able to sit through a family dinner without checking text messages, but some people like to check them. Some people check their emails every five minutes but other people don't. The research is very interesting, and it is certainly something we are taking on board, but I do not think it is our position to say that, because it depends very much on the individual.

Q534 **Ian C. Lucas:** I think it is exactly your position to say that because it is your product. Young people and children use Twitter. I use Twitter. I think I am addicted. My wife thinks I'm addicted. One reason why you are so effective is because people are addicted. The analogy is smoking. That does not mean that we can't regulate and take appropriate steps to manage it. My question is whether you agree that Twitter is addictive?

Nick Pickles: In a personal capacity, I am able to go several days without tweeting.

Q535 **Ian C. Lucas:** How many days in a year?

Nick Pickles: It is public, so you can check. It is an incredibly important question. We are working with the Secretary of State for Health, who has established a working group on this, and we are working with the Department for Digital, Culture, Media and Sport on the digital charter and the internet safety strategy. These questions are very important. They tie into digital literacy and how we equip young people to understand how the world is around them and how to make assessments of the veracity of information. For those sort of questions, I think that is where academics are best placed. We in the industry have to work with them to make sure—

Q536 **Ian C. Lucas:** Academics have told me that it is addictive.

Carlos Monje: I spent a portion of yesterday at an event run by Common Sense Media, which is looking at exactly these questions. They ran a very entertaining campaign that featured the comedian Will Ferrell talking about the challenges of having a device-free dinner. I am glad that you like the product; I hope you are active on it right now. These are really

important issues and we think about them very carefully, especially for our younger users. We don't allow folks younger than 13.

Q537 Brendan O'Hara: We heard from Google and Facebook this morning, and both said that the trust of their user is one of the key elements to their business. Is the trust of your user as important to you as Facebook and Google claim it is to them?

Carlos Monje: Definitely. People come to Twitter to figure out what is happening in the world and to talk about what is happening in the world. If they have a bad experience, they will find somewhere else to go. We are very proud and very jealous of our 330 million users. We want to keep them. We know that there are other places that they can go and we know that we have competitors that are trying to take them away. Part of that is giving them a sense of safety and confidence in what they are seeing.

Q538 Brendan O'Hara: The reason I ask that question is because, in your answer to Mr Hart, you seemed to say it is like the wild west out there—an absolute free-for-all—and you won't act on what is a lie or anything that is out there. How can you build trust among users on a platform when you yourself say that there are no circumstances, other than legal, I presume, in which you would be arbiter?

Nick Pickles: I think you are familiar with the debate about whether Scotland would have had access to the pound after the independence referendum, and we are all familiar with the debate around whether Britain sends £350 million to the European Union. One of the strongest characteristics of Twitter is that our users are the very people who are engaging in those debates, challenging people who share information they disagree with and sharing their version of events as they see them. Our job is to make sure—this is where the Committee's work is important—that those actors who seek to manipulate our platform are not able to manipulate, and that credible voices such as journalists and the media partners we work with, are elevated. This is an important question because it goes to a lot of issues in the industry and also internationally with parliamentarians. I do not think people want technologies to be telling them, "Sorry, you've said something that's untrue. We're going to delete your account now." It is about elevating credible voices so our users can investigate issues, speak to other users and come to their own conclusions.

Q539 Brendan O'Hara: You talk about the manipulation of the platform. The platform is generally regarded to be being manipulated. How concerned are you that Twitter seems to be the No. 1 platform of choice for those who seek to disseminate fake news or disinformation? Why would Twitter be identified as the weakest link in the chain for those who wish to do that?

Nick Pickles: There has been some really interesting research in the past few weeks about the reach of fake news and junk news. Oxford University published a study. I challenge your premise. We were really excited to see recent numbers showing that the number of people visiting news publishers from Twitter is significantly increasing. That is a partnership we



HOUSE OF COMMONS

really value. There is a huge number of journalists in parts of the world—journalists like the guys at Bellingcat, an organisation that did not exist a few years ago but is doing groundbreaking verification research—

Q540 **Brendan O'Hara:** I have no doubt at all that that is true, but on a platform of 330 million users, I would expect the vast majority to be, first, human and, secondly, truthful, honest and decent. You are seen as the weakest link in the social media chain for the dissemination of fake news and false information. I am wondering why.

Nick Pickles: That is probably a conversation to have with someone like Claire Wardle, who is on the next panel. She would probably say that the term “fake news” in itself is unhelpful in understanding this problem. We have a lot of politicians around the world who use Twitter. We have a lot of journalists who use Twitter. A lot of different perspectives are shared. There are also people sharing tweets from the moment something happens in a country. Often, the first time we see it has happened is from a photo or a tweet from someone on the ground who often does not realise what is happening. There is a huge amount of information on Twitter, and our job is to surface the credible information to help our users.

Carlos Monje: May I further challenge your premise? One of the things that Nick mentioned earlier that I think it is important to go back to is that researchers tend to flock around our data because ours is the only data out there. Maura Conway, who is a professor at Dublin City University, talks about the over-indexing of the challenge to Twitter. We are not denying that people try to manipulate our platform. We are getting better and better at stopping it. This is not a challenge that is going to solve itself quickly. The folks you have asked to testify in front of you today provide a good cross-section to try to understand that. We recognise that this is a challenge and we work every day to make it a better experience. Part of that is the kind of progress that we have had: 6.4 million accounts a week. That is up 60% just since October last year. That is 3 times where it was just two years ago. We measure our progress in weeks and months. Like the folks who testified earlier, our bottom line is tied to getting this right.

Q541 **Brendan O'Hara:** I accept that, but we have heard an awful lot during our investigation and our discussions about—it is an old rule, but it is never usually wrong—following the money. Having many millions of fake accounts or bots among those 330 million users allows you to sell advertising and claim a potential reach to advertisers far beyond what you could without the fake accounts. In what way have you become a prisoner of the bots?

Nick Pickles: We do not include those spam accounts in our metrics that we report to the market.

Q542 **Brendan O'Hara:** How many bot accounts are there?

Nick Pickles: We can share that information—it was shared with the stock market this morning. It is less than 5% of our total user base.

Q543 **Brendan O'Hara:** 5% of 330 million?

Carlos Monje: No, the 330 million figure excludes the 5%. To your point, I think that the exact opposite is true. If people do not have faith in what we are saying, they will not be willing to come to the platform or to advertise on the platform. That is why we fight this with all the tools that we have.

Q544 **Brendan O'Hara:** Do you accept that every time you identify a fake account or a bot, there is someone, somewhere, in a bot farm churning out thousands more? You are fighting against the tide here. Are you absolutely convinced, and can you convince us, that you are doing absolutely everything that you can to combat the bots?

Nick Pickles: There are two things there. First, the benefit of Twitter for academics is the open API. That allows people to build apps, which are often part of that bot problem. We have suspended 155,000 of those applications since the middle of last year, which were collectively responsible for 2.2 billion low-quality tweets.

One thing I would say is that being a bot per se is not against our rules. There are plenty of bots around the world—some of them edit Wikipedia, some of them monitor air quality, and some of them tweet out images from web cams in real time. The idea that an account is automated, and is a bot, to use the colloquial phrase, is not in itself a bad thing.

Something that is concerning is that last year, if you saw something you did not like, you said, "Oh, it's fake news." That was how you discredited a source. This year, we are seeing the same thing happen with people replying to real people on Twitter and saying, "Oh, it's just another bot," as a way of disparaging their viewpoints. That is not helpful. I think it is important to distinguish between malicious automation and bots that are actually serving a useful purpose on the platform, many of which are used by news publishers to share their content in breaking news scenarios.

Q545 **Brendan O'Hara:** Yes, but in the great scheme of things, how many of the 6 million or 7 million that you have identified check air quality, and how many would be deemed to be malicious? You are not looking at a 50:50 split, are you?

Nick Pickles: We are removing a lot of those accounts. One of the challenges that we have—we invest a lot in this technology—is stopping somebody coming back when we remove them.

Q546 **Brendan O'Hara:** When it is done on a factory basis, and they just keep coming and coming, can you be confident, and assure your users, that you are doing absolutely everything that you can to stop the factory production of fake accounts?

Nick Pickles: Absolutely. We have a dedicated information quality team in the company, which brings together people from our news team, our engineering teams and our safety teams to make sure that that problem is being addressed. It is also why the industry partnership is important. The Committee has heard about the advertising model, and how many steps



HOUSE OF COMMONS

there are between the user and the person who is the advertiser. Infrastructure on the internet works in the same way; there are lots of different partners. We, as an industry, are doing more together, because, to use your analogy of a server farm, they are getting SIM cards from somewhere. How are they using the phone systems to create the accounts? We need to make sure that there is a top-to-bottom response across the industry, not just at the surface layer.

Q547 Chair: Just to be clear, 5% of a user base of about 330 million would be roughly about 16 million bot accounts?

Carlos Monje: Those are rolling averages. The 330 million is not our total user base; it is the people who check in on a monthly basis. We are constantly knocking bad accounts off our platform. They keep trying to get back. We are getting better at stopping those reincarnated accounts by looking at signals, where they are dialling in from, and the device information to the extent that we have it.

That has been the key to success for us, by the way, in fighting terrorism—more than 90% are taken down before anybody else reports them; 75% before they can tweet. The way we do that is getting better and better at informing our AI—our machine learning. Our teams are industry-leading on some of these challenges.

Q548 Chair: What was the figure you said? How many accounts every day does Twitter take down?

Carlos Monje: We don't disclose that.

Nick Pickles: We can share the information that was disclosed to the market this morning, which has our most recent numbers, and give you a breakdown of all the takedown numbers as well, if that is useful afterwards.

Q549 Chair: But do you know what the takedown number is?

Nick Pickles: I am just trying to double-check. 6.4 million accounts are challenged every week.

Carlos Monje: We just announced new numbers this morning. We will have to come back on that.

Q550 Chair: Okay. So if it is that number a week, you are basically challenging, every year, the equivalent of the entire active user base on Twitter. That would suggest that there is a massive problem.

Nick Pickles: It would also suggest that the steps we are taking are showing results in detecting activity and stopping it, which is the important metric for us.

Q551 Giles Watling: I was going to come in on the subject of bots, but we seem to have covered that subject. I am still sitting here in absolute shock about the responsibility that you take or do not take. We are not talking about an 18th-century broadsheet that is tomorrow's chip paper; we are talking about something that reaches into every corner of society.



HOUSE OF COMMONS

We are talking about a Twitter storm—those famous words—that takes off and then becomes the truth because it is repeated often enough and people start to believe it. I am still astounded.

With this enormous reach you have, you have enormous power, and with enormous power comes great responsibility. You seem to want to duck that and say, “Well, it’s not up to us to decide whether it’s true or not,” but I suggest to you that you have to start taking that responsibility on board. A very small percentage shift in people’s views, which you can do through your Twitter accounts, can change the course of an election or a referendum, and it might well have done—we don’t know. I believe that you have to start taking that kind of responsibility. Is that something that you are ever going to take on board or is it something that you are just going to carry on ignoring?

Nick Pickles: To clarify, the responsibility that you are talking about is removing content based on whether it is truthful or not.

Giles Watling: Yes.

Nick Pickles: During an election campaign in the UK, political advertisements are exempt from the advertising rules, so that would be taking regulation of UK political advertising and giving it to American technology companies. In terms of the democratic process, that seems to me quite a robust step to take, and it is not one I think we want to take.

Q552 **Giles Watling:** I absolutely take your point. What I am saying is: do you not have some sort of moral compass as a company that says, “We’ve got to start addressing this”? The routes you might have to take to address it might be difficult—you might have to deal with issues of law and so on—but is it not something that you would even begin to think about?

Nick Pickles: That is why I answered the Chair’s question earlier by saying that context is essential. We do see people who create impersonation accounts, where they pretend to be another person and use that as a way of destroying their reputation. We may also see someone create an account that uses hateful language, as we heard earlier on, to disparage someone based on their gender or a protected characteristic. That responsibility is absolutely one that we take very seriously, because that is the behaviour of our users.

I think there is an important distinction between the behaviour that our users exhibit and the content of the ideas and viewpoints that they hold. That is a big public debate that we have not started to have yet: where does that responsibility lie? You heard from previous witnesses that in Germany a debate is now happening about whether they have found the right balance. That debate is one that is very, very important, and this inquiry is a very valuable opportunity to start having that discussion. You have some great witnesses this afternoon who will probably have strong views on that as well.

Q553 **Giles Watling:** I am looking forward to seeing you begin to push that



HOUSE OF COMMONS

argument along. Do you not feel that you might risk losing trust if Twitter is seen to be disseminating fake news and not taking it down or dealing with it? It becomes an existential problem for Twitter itself: if you lose that trust, you lose the business and you are out of business. You said yourself that there are people trying to move in on your territory now. Is that not something that you are facing now?

Nick Pickles: When you look at the problem of what has become colloquially known as fake news and break it down, one of the biggest challenges is hyper-partisan content—content that is true in some ways, but is very aggressively spun to match a politically partisan viewpoint. The research from Oxford—and from Dartmouth, I think—found that consumption of that also goes along partisan lines. There is a partisan community feeding themselves on content which furthers their own partisan views. I think we take very seriously as a company our role in helping people see broader perspectives than the ones they hold. To go to the earlier point about how we manage information, we do not want to intervene in that process and push people more into those communities of hyper-partisanship, but we are very interested—we are discussing this with academics—in how can we be a positive value to help give people access to a broader range of credible information, from credible sources. Projects like the trust project, mentioned earlier on, are incredibly valuable to how we come across these problems, but there is a delicate line to tread, because in some countries the Governments would very much like us to remove political opponents and independent media. We have to ensure that, as a global platform, we do not take a step in one country which would seriously harm civic debate in another.

Q554 **Giles Watling:** Do you not feel, to a certain extent—I mentioned this earlier to other people—that you are a step behind everything? These might not be intended consequences of your business, but out of them come these extraordinary events, like the Twitter storms we were talking about. Do you not accept that perhaps you are one step behind and you have got to get ahead of this game? Otherwise, you will lose confidence and trust, and ultimately you will lose business.

Nick Pickles: As I say, the company right now is both developing new technology and deploying technology from other purposes that we can reuse. You may have seen that we paused our verification process. One of the questions at the heart of pausing that process was how we can leverage that process to better inform our users about information. So yes, we do not want just to fix today's problem; bad actors will change their methods, and we want to make sure that we are keeping ahead of the problem.

Giles Watling: That, at least, is good to hear. Thank you.

Q555 **Jo Stevens:** Mr Pickles, I think it was you who mentioned earlier Britain First: a fascist organisation whose videos President Trump retweeted late last year. I am sure you will recall that part of the evidence in the trial of the man who assassinated our friend Jo Cox MP was that he shouted "Britain first" before he killed her. Leaving aside the bots issue, I am



HOUSE OF COMMONS

really concerned about the flaws in the policing system for your platform, because it was only after President Trump re-tweeted those videos from Britain First that you took Britain First down from the platform. How can your system be described as anything other than inadequate?

Nick Pickles: I am grateful for the opportunity to discuss this, because it was a policy question we discussed a lot internally. The actions of any user in re-tweeting another user are not factors in our policy framework. We had announced that we were reviewing our approach to non-terrorist violent extremist groups. Our policy prohibits the promotion of terrorism, but there were certainly groups that fell into the non-terrorist organisation space. We then took the best part of three months to try to understand where the lines of that policy would fall, not just in the UK but, as has been mentioned, in India, the US and other countries. When we enforced that policy, Britain First was one of the organisations affected, but it was not a decision to remove an account because of another account's actions.

Jo Stevens: It was just a coincidence.

Nick Pickles: It was. Actually, we took accounts down around the world relating to that policy change, because, as I say, we think that, as we change our rules, groups change their behaviour to try to stay on the right side of our rules. So we felt that was an opportunity to take action on accounts that we felt should not be on the platform.

Q556 **Jo Stevens:** So I should not read anything more into it than a coincidence of timing.

Nick Pickles: Yes.

Chair: A quick question from Rebecca Pow, and then Paul Farrelly will finish off.

Q557 **Rebecca Pow:** I was staggered, like Giles, to hear you say, "We are not the arbiters of truth. We do not have rules on truth." To me, that gets to the nub of what we are all talking about today. As a platform, you are openly able to spread disinformation, which is tweeted and re-tweeted and shared. What worries me is what this is doing to our children. Shouldn't you take some responsibility for it? If you cannot and are not able to and your policing system is not up to it, surely some sort of regulation or body will have to be put in place to ensure that the next generation is safe.

Carlos Monje: That is an excellent question, ma'am; thank you. I do not know if in the UK you have the same experience of conspiracy theorists and folks you cannot convince they are wrong about a particular issue. We have seen our peer companies try different things that have actually been counterproductive and where fact-checked articles get more attention than they would otherwise. We feel the best way to have a healthy platform is to be very good at what we are good at and to be humble about what we are not. Malicious use of automated tools to manipulate our platform, is something we can stop. It is against our rules.



HOUSE OF COMMONS

When it comes to elevating credible voices, I will give you the example of the Boston marathon bombing a few years ago. It was a terrifying moment for the people in that city and for a period of time people were locked in their houses. Then a story went out that the suspect had been captured, and it was not true. One of the things we do as part of our philanthropic mission is to work with emergency responders across the globe to make sure that they can weaponise our platform to get the truth out more quickly. It was actually the Boston Police Department that said, "No, we don't have a suspect in custody. Stay in your homes." We feel like that is a more appropriate place for a platform of our size to be.

The other piece—

Q558 **Rebecca Pow:** That is a really good example, and thank you for sharing it, but it seems completely hypocritical as well that on the one hand you can say, "We don't arbitrate the truth", and on the other hand, you are running "trust courses" and things. It seems like a muddle.

Should you at least have a much better rating system, so that when people see tweets they know whether it is fact or fiction, so that they can at least have some sort of gauge about how to rate this in their minds? And how on earth are children going to manage to deal with this?

Nick Pickles: There were two questions. I am not familiar with the training you mentioned in Whitehall.

Q559 **Rebecca Pow:** No. You said you run a trust programme.

Nick Pickles: Sorry. There is a third-party organisation called The Trust Project, which a number of companies—a number of media organisations—are using, to bring—

Q560 **Rebecca Pow:** Yes, but I am asking you about a rating system. This has been mooted to us by various academics and people we have met, which would include some sort of register of whatever they are seeing—whether it is fact, fiction, half-fact, half-fiction, or something like that. Is that possible?

Nick Pickles: There is an interesting question there, in terms of, "How do you verify source?" Do you tell someone, for example, "This account is very new", versus, "This account has been around a long time"? "This account is satire" versus "This account is not satire"? One of the things that researchers have flagged up is that the danger of those kinds of labels is that they do not work if the person's cognitive bias is inclined to agree with the subject of the story. So, even if you put that label on something, if it is about immigration being too high and you think immigration is too high, you jump to the subject of the story, rather than reading the label.

That is why The Trust Project, and the academic research we will be funding and partnering with, is coming up with solutions to that problem. Absolutely, we think there is more information we can give to our users. The question of what information is useful is the incredibly difficult one

right now. Our verification scheme—actually, review—is looking at these questions as well.

Rebecca Pow: I just put it to you, Mr Chairman, that more academics would like to join with you on coming up with the solutions.

Q561 **Paul Farrelly:** Mr Pickles, I am still concerned about your categorical statement about the case of a Conservative candidate in an election where lies are demonstrably told but you do not see it as your role to be the arbiter in any way. I am concerned particularly because lots of people out there do not have as thick a skin as politicians grow, and lies that are disseminated can lead people to commit suicide, for instance.

Locally to me, we have an anonymous malicious Twitter account that spreads innuendo about me and other people. It does not affect me because I never look at it, but it affects other people and their mental health. I think we know who is behind it and I take the view, “Don’t look at it, then”, because this person is a sad, miserable, pathetic excuse for a human being. But I do have some concern about the effect it has on other people who do look at it.

So, if I came to you, say, with lawyers and asked you for details of the person—assuming you knew, or not—or the IP address of the person who had established that account, what details would you give me? And I am not President Assad. I am an ordinary Member of Parliament doing a decent job, I think.

Nick Pickles: Obviously, you will appreciate that that data would be covered by data protection law, so we would be required to ask for legal process to allow us to disclose that information. The police, for example, seek that legal process on a daily basis, and we assist in thousands of investigations every year by providing user information, but we would require legal process to produce that information. I am happy to make sure our teams have reviewed the account you mentioned afterwards.

Q562 **Paul Farrelly:** So would I have to go to court to get an order?

Nick Pickles: Yes.

Q563 **Paul Farrelly:** Rather than you co-operating? What if I cannot afford to go to court to get an order?

Nick Pickles: You raise a very valuable question about access to justice. Parliament recently passed a law to change the process on defamation to address that very question. If there is a concern about accessibility to the justice system, parliamentarians are best equipped to address that.

Q564 **Paul Farrelly:** But you can see that, if you are warned that something is a lie and if you are warned that it is having severe, possibly life-disturbing effects on people who cannot afford to go to the law, some right-thinking people might take the view that you are complicit and therefore should face a liability.

Nick Pickles: There is an existing legal framework for that process, and Parliament recently amended it to improve it. That legal process involves a



HOUSE OF COMMONS

third party—not companies—deciding on the veracity or not of the conversation at hand.

Q565 **Chair:** Final question. I have had cases of people coming to me very disturbed that their identity has been stolen and they are being impersonated on Twitter. To my knowledge, they have not got anywhere with you when making complaints. What is your policy, if I were to come to you and say, "Someone is impersonating me; I would like you to remove that account"?

Nick Pickles: We can send you the full text of the policy. We do not allow impersonation accounts. We do allow parody accounts, but those accounts have to very clearly state that they are parodies. You may be familiar with the @Queen_UK account or the @PeterMannionMP account. They're not really the Queen or Peter Mannion—

Q566 **Paul Farrelly:** I don't follow Twitter. Life is too short, I'm afraid.

Nick Pickles: As I say, we have a policy on impersonation and we remove accounts that break it. I am happy to follow up on the cases.

Q567 **Chair:** A final question for Carlos. We heard earlier on that Twitter challenges more than 6 million accounts every week. How does a company that employs fewer than 4,000 people challenge 6 million accounts a week?

Carlos Monje: That is an excellent question. For us, the number to keep in mind is that there are 500 million tweets a day, and there are 350,000 tweets a minute. The key for us is to use machine learning.

We acquired an incredible British company called Magic Pony, which has brought not only technology but incredibly talented engineers and data scientists. They help us train the machines to make consistent decisions about what is and is not automated. Some of the things we look at are near-instantaneous re-tweets. If Nick tweets something and a fraction of a second later I re-tweet it, that is not human behaviour. It takes a moment to open the app, to read the information and to decide that I want to tweet it. A certain regularity in the timing of tweets is, again, not human. Sometimes people in a moment of excitement will tweet a large number of tweets in a given time period. What we will do, based on the certainty that we have using our machine learning, is to challenge those accounts. So you are right, it is not people all the time. We have a combination of human review and machines that help us police our platform.

Q568 **Chair:** Thank you very much. Thank you both for your evidence today. That concludes the evidence scheduled for this morning. We were due to restart at 1 pm Washington time. Given that it is already three minutes past 1, we have probably missed that target. I suggest we reconvene at about 1.45 pm.

Examination of witnesses

Witnesses: David Carroll, Associate Professor of Media Design, The New School; Amy Mitchell, Director of Journalism Research, Pew Research Centre; Frank Sesno, Director, School of Media and Public Affairs, George Washington University; and Claire Wardle, Research Fellow, Shorenstein Centre on Media, Politics and Public Policy.

Q569 Chair: Good afternoon. Welcome to this session, which is part of the Committee's inquiry into disinformation and fake news. We heard in the previous session from the tech companies, and we are delighted to be joined by people who are leading the investigation and research into the issues we were discussing in the first session. The Committee have direct questions for you, but as the session goes on we would certainly welcome any immediate observations you have on the evidence that we received from the tech companies.

I wonder whether I may start with David Carroll. In the session with Facebook, a reference was made to an investigation that the UK Information Commissioner's Office is conducting that is linked to Cambridge Analytica. This is your case, so I thought perhaps you could tell us something about it. My understanding is that you made an application to Cambridge Analytica for data that they hold about you linked to the election period in America in 2016, and that because at some point that data must have been within the UK jurisdiction, the UK Information Commissioner's Office is looking into that. This is a really interesting area. I am sure that we will be very interested in following this up with the Information Commissioner when we are back in the UK, but we would welcome any insights you can give us into your case and what you are trying to achieve.

David Carroll: Sure. Thank you very much for the opportunity to address the Committee and to share information about this transatlantic investigation.

Back in February 2017, I was advised that I could make a subject access request under the UK Data Protection Act to the company Cambridge Analytica because of the reports that it had been hired by various campaigns during the 2016 campaign season. Someone who has data that is being processed has right to have it provided to them. I received the data from the parent company, SCL Group, at the end of March. It was accompanied by a letter indicating that they were trying to comply with the UK Data Protection Act, some description of how the data was collected but nothing specific, some description of the kind of third-party entities that the data would be shared with but nothing specific, and an Excel spreadsheet that contained accurate voter registration information, a tab of election returns relevant to my voting district and a panel called "models", which was an ideological model that tried to predict and analyse my political beliefs. It comprised 10 political topics, which were ranked in order. It also tried to compute my partisanship, separate from my registered partisanship, and tried to compute my propensity to participate.



HOUSE OF COMMONS

I decided to post this on Twitter on the day that I received it, while redacting my personal information. I was soon able to get a solicitor to represent me in the UK to challenge compliance with the UK Data Protection Act, because we believed that the disclosure was not complete and might not be compliant. At the same time, I was advised to file with the Information Commissioner's Office, which I did on 4 July 2017, which is a little bit poetic. I was able to contribute to its ongoing investigation into voter analytics companies with reference to the elections in the 2016 season. So there is at least one US citizen who has filed with the ICO.

We have received updates recently that the investigation is far more complex than anticipated. The ICO had hoped to provide a report last fall, but as they dug deeper into it, it became more and more complex, so they are still working on it. But we have received follow-up that the investigation is coming along. It certainly has been useful to provide this information also to the Senate Select Committee on Intelligence and others here in the US who are trying to figure out the company's role in the 2016 election. The key idea is that US voter data appears to have been processed in the United Kingdom, and this is probably unprecedented.

Q570 Chair: Why do you think that has happened? Is it simply because of the construct of the company, Cambridge Analytica and its groups, that that data has been processed outside America? Are there any rules around where data is stored and processed relating to politics in America?

David Carroll: Interestingly, it does not seem like US law was potentially violated here. It really exposes how there are not sufficient US protections for voter data. It also exposes how US citizens do not have the same rights that UK citizens and EU citizens enjoy—the ability to get data controllers to disclose data held on them when they become a data subject. This experience has illuminated a lot of things, but one of the things that is most stark for us is that the US has inadequate protections in this regard.

Q571 Chair: This is data that Cambridge Analytica is holding on you. Do you know where that data was acquired from, or do you hope to be able to discern that when you get access to it?

David Carroll: There was no indication of where they obtained the data, and that is part of the insufficient disclosure. We should be able to know where they got the data, how they processed it, what they used it for, who they shared it with and also whether we have a right to opt out of it and have them delete the data and stop processing it in the future.

Q572 Chair: You heard what Facebook said in the earlier sessions about the right to request data deletion—whether or not you cease your Facebook account—and the fact that Facebook developers are required to give up any data that they acquire from Facebook users who request that. Do you think that, in reality, it is as straightforward as that?

David Carroll: I have myself downloaded my Facebook data, as they described, and it is certainly by no means a complete disclosure of all the



HOUSE OF COMMONS

data that Facebook has on a user. But I also know that the company has pledged to try to be compliant with the general data protection regulations and will be launching new controls and new disclosures for users to try to be compliant with that new privacy regime. I look forward to seeing whether users are able actually to get a complete profile when they download their data in the near future.

Q573 Chair: I should like to ask Claire Wardle a question before bringing in colleagues. You will have heard what was said in the previous panel, particularly what Twitter said about their policies on not just fake information but lies. Do you think that it is irresponsible of them to have that attitude? Surely, if a major platform like Twitter believe that they have no obligation to take down disinformation, that is a severe inhibitor of our efforts to combat the spread of issues such as disinformation that we would all regard as being a social harm.

Claire Wardle: This really does get to the root of the question, and it is obviously what you were saying. It was interesting sitting behind as some American colleagues were saying, "First amendment, first amendment!" That has not come up yet in this discussion. The difficulty here is that most of this information is not true or false; it is somewhere on a spectrum. We take specific examples and we want to say that that seems wrong, and we have to be able to take that down. The truth is that the scale of this makes it much harder. Many people believe that there is not an obvious truce. Where we know that something is 100% inaccurate, how do we talk about that? This gets to the question of definitions. When we are talking about this huge spectrum, we cannot start thinking about regulation, and we cannot start talking about interventions, if we are not clear about what we mean. What is misleading and what is hyper-partisan? As a British person, I could argue that much of the content of our press that is published every day is misleading in some way. Whilst it is hard to hear Twitter say what they are saying, because from a particular example, and we want to say that that seems wrong, I think that we would be in very dangerous territory to start saying that something clearly needs to be taken down, and something else does not. Who would do that? I wish that we did not live in that world, but we do. We go to the pub and we all say different things and we gossip. That is partly what makes humans humans, and it makes us feel uncomfortable, but I cannot imagine where we would even start if we said that we lived in a world where we could make those decisions, and I do not want them to make those decisions.

Q574 Chair: But there are times when there are demonstrable lies. In the debate about disinformation and fake news, it is quite clear, as you said—your analysis and academic work demonstrates the different scales and gradations of fake news—that there are, at the most extreme end, things that we know are a lie. Surely it is particularly dangerous not only for that sort of information to spread uncorrected, but for it to be spread anonymously as well. The courts are open. In fact, UK people have faced legal redress because of Tweets they have shared or posted on Twitter, because we know who they are. It has got to be even more concerning



HOUSE OF COMMONS

when people use the protection of anonymity to spread lies about other people.

Claire Wardle: No, and it makes us feel very uncomfortable. Because Twitter do not have a real name policy, it makes us feel very uncomfortable. But at the same time, there are many people who have good reasons for not using their real name. We have to think this through. Twitter talk about themselves as a self-cleaning oven; they argue that many journalists would correct lies very quickly. There is no way to tag that correction to the original problematic piece of content. That is what I would like to see. When journalists, civil society and fact checkers work on that, how can they not flag it up and say that it is false but connect alternative pieces of information to it? So if we have a healthy debate about what is true or false, that is around that original piece of content, because at the moment they are completely separated and that is something that we need to look at from a technical point of view.

Q575 **Simon Hart:** Can I have some clarity on one point that you have just made—for which, thanks. You said, “I dread to think where we may start on this.” When we were talking amongst ourselves over lunch, we said that where we start is where we already are, as far as print and broadcast media are concerned. There are already established norms for people who have the power to affect the outcome of elections by what they choose to print and what they choose to withhold. I suppose that we are asking whether there is a sustainable argument that explains why people who run an online platform consider themselves to be in a very different place legally from those who run an offline platform such as a newspaper. If there is, we have not heard it yet.

Claire Wardle: I completely agree with you. My frustration is that we get into these battles of definitions, when we say, “You’re a publisher” and the platforms say, “No, we’re a platform.” The truth is that they are somewhere in the middle. They are a hybrid form of communication. What I would like to hear—to be honest, I did hear some of it this morning—is, “We would like to be part of a conversation about what new forms of regulation might look like.” I do not think we can use the broadcast model. We cannot regulate speech on Twitter in the same way that we regulate the BBC. That is not workable, but we should not say that there should be nothing in that space. My frustration is that, in all these conversations, we are not actually getting to what this new hybrid form of regulation looks like. That is where we need to get, and we need to get there quite quickly.

Q576 **Simon Hart:** I have a quick follow-up question. Some of the evidence that we have heard this week, from tech professionals and from academics, is that the moment for self-regulation has probably passed. The social media platforms do not appear to be taking it seriously—certainly not as seriously as they take their commercial objectives. Therefore, the only option left is some kind of state intervention, maybe light touch. Do you think that the social media companies have cottoned on to the fact that this net is now closing around them?



HOUSE OF COMMONS

Claire Wardle: They have absolutely cottoned on to that. But I would like them to come out and say, "Here are some possible solutions that we are all behind. Can we start this conversation?", instead of sitting in the trenches, saying, "We are not publishers." I sit on the European Commission's high-level group on fake news and online disinformation, and this approach is coming. The conversations in Europe are very different from the conversations that you hear in the US. But I do not think that we should have state intervention that may be a knee-jerk response and does not react to the realities and challenges that come from these platforms at a scale that is hard even to imagine. I want the platforms to be part of the conversation, so that we can have an honest look at the issue. We never saw this coming. We did not think that, in 2018, we would be where we are today. We should not be the only ones who make these decisions and they should not be the only ones.

Q577 **Simon Hart:** That is all good stuff. I agree. But we did not see much evidence of that this morning. Surely this morning was the ideal opportunity for us to get at least an indication that that thought process was happening.

Claire Wardle: Yes. The social media platforms said that they wanted to be part of the conversations, but I did not see them coming up with potential solutions.

Chair: Ian, do you have a quick question on this?

Q578 **Ian C. Lucas:** My question relates more to the line that I was pursuing this morning on electoral issues and regulation. I was very interested in what the representatives of Facebook, in particular, said about the different rules that they are beginning to introduce. They kind of conceded that past elections had shown that there was a capacity for the law to be broken as a result of failure to disclose information. What was your reaction to the proposals that they made and how do they sit with, for example, US electoral law?

Claire Wardle: My concern is with definition. If we just talk specifically about elections, are we just talking about what candidates or campaigns push out? If we look at some of the evidence we have of Russian interference in 2016, it was about posts that had nothing to do with obvious political issues. They were cultural and social issues.

My concern is having an honest policies ad that says that all the candidates and campaigns have to be transparent, but that does not look at all the stuff that is actually causing problems. My fear is that we create boundaries that actually do nothing because the real problem is outside those boundaries. We are talking about elections and political content, but look at something like Pinterest and the visuals about anti-vaccination information. This is not just about politics. We need to be a little bit careful about that.

I was a little bit surprised that the platforms admitted that, in terms of UK electoral law, they were not looking for things that it seemed pretty shocking not to look for. They were honest that people can lie about their



HOUSE OF COMMONS

location, but the fact is that they did not think this through earlier. It was surprising to me that they did not think, "Isn't this an issue if people are buying advertising and we cannot independently verify their location?"

Q579 Ian C. Lucas: That's obvious with Twitter in particular because it does not have an address. That can really have an impact, particularly in view of the volume and reach of Twitter. In the UK, there are national rules on no finance from outside the UK and local spend is very important. This makes it even more difficult.

Claire Wardle: But even on the most basic level, I could post now to Facebook and it would say, "You're in Washington DC". But I could override that and say, "No, I'm in Antarctica." That is the same for every social platform. The fact is that I can just override the system and make up my location. There seem to be basic location issues on the platform that we should think much more carefully about. Of course, there are privacy implications, but the idea that I can override my location is problematic.

Q580 Ian C. Lucas: What about access to the data that the platforms have? They seem to think that the fact that they were looking at it and that they would put the information out was quite sufficient. My question to them was, how is it policed? How do we know that they are actually complying with their own rules? We need some independence.

Claire Wardle: No, absolutely. I would hope that they would recognise that they need independent auditing of not just what is on the platform now but the steps that they are taking. So there should be independent auditing of the data for the new "fact check" tag that Google has, for example. There are ways in which they can make that data anonymous, so they should have ways that we can sample data and we can sample content, in the same way as public broadcasters in Europe commission independent civil society organisations to audit the output. There should be a way in which we can do exactly the same thing on the social platforms. That seems to me to be a kind of low-hanging fruit for them—it is in their interests to have people to look. At the moment, all the conversations we are having are because excellent journalists like Craig Silverman at BuzzFeed and Julia Angwin at ProPublica have started them. Some of the examples you gave today are about journalists going to the platforms and searching, and they have limited access to the data. Imagine what we could really find if we had true access—that is why they do not want to give it up. For me, we are past the point of believing that they are going to give us that information ourselves. I would argue we have to have access to that data.

Q581 Jo Stevens: Claire Wardle, I would like to go back to some of the points I raised this morning about algorithm secrecy and inherent biases, as against transparency, trust and the users' understanding of what is happening when they go on platforms. What do you think the solution is? Where can the balance be found between those two competing interests? Is there an obvious or simple solution?



HOUSE OF COMMONS

Claire Wardle: I thought your question got to the heart of so many things that we are trying to discuss today. Most users do not understand this space. If we are talking about news and media literacy curricula, that has to include teaching about how to evaluate an algorithm and how to understand how what you see on Amazon, Netflix or Facebook has been decided by an algorithm, how an algorithm gets developed, how it is created by a certain person and how their biases might shape that. That has to be part of the teaching that we give to people.

In terms of the platforms themselves, there has been some excellent work on algorithmic transparency. The Tow Center at Columbia—I think in 2014—talked about algorithmic transparency and what a framework for it might look like. What are the questions to ask a platform about why it was created? What are the metrics for that particular algorithm? How can we have more insight into that algorithm? How can we think about frameworks of algorithms? Irrespective of the platform, how can we set up that framework so that platforms have to be not just transparent but transparent across particular aspects and elements? I think that is the key. We keep talking about algorithms as black boxes. We need to start saying, “How do we get into those black boxes in a consistent way?”

Q582 **Jo Stevens:** Does that not go against the whole ethos of these companies? Their entire business model is based on secrecy, isn't it?

Claire Wardle: They would talk about the secret source, and we would understand that there is a part of that, but when their secret source affects the information that people receive, there are ways of coming up with a framework that is not really about them sharing their competitive edge. It is saying, “Can we talk about why that algorithm was designed in the first place and what the metrics are?” If we were Google, we would say, “To surface the highest quality information as quickly as possible.” I think we need to be asking those questions. Yes, they are commercial companies and we have to understand that, but they have influence because they have become the dominant source of information globally and they have to understand that responsibility. All three companies said, “Yes, we see this as one of our responsibilities.”

Q583 **Julian Knight:** Mr Sesno, what do you think the likes of Google, Facebook and Twitter fear from policy makers, not just in the US? We understand that the US is caught up in quite a bipartisan argument about the concept of fake news and its impact in the presidential elections. What do you think it is that these companies fear? What do you think would be the most effective course of action to get the acute problem of disinformation more under control?

Frank Sesno: What do they fear from you and our elected representatives? Regulation. They fear that they will not be what they profess to be, which is technology companies and not media companies. They fear that they are going to be held to account for the content that they say they are merely facilitating and not producing. The most poignant observation is that they have this very strange, powerful, hybrid identity



HOUSE OF COMMONS

as media companies that do not create any of the content but should and must—to their own inadequate levels—accept some responsibility for promulgating it. What they fear most is regulation—a requirement to turn over their data. They fear that there will be Government regulators at different levels overseeing their businesses and that they will not be able to be the independent mega-corporations with the mega-revenue that they now generate.

Q584 Julian Knight: What do you think is the most effective means by which to get some control?

Frank Sesno: Well, conversations like this are very effective. They are feeling the heat in a very powerful way. They are also feeling the heat in a very powerful way from the journalistic community—from the publishing community. I thought it was very interesting to hear them talking about how much they are working with publishers to help them generate revenue. Well, talk to the publishers. What is very powerful and very prevalent now is to make this conversation as stark as it is and to put online what is on the line, which is whether we are going to have an informed or deformed public discourse and public process and whether those companies are contributing to or subtracting from the democratic health that we value. I do not think that we can paint this in stark enough terms. We have to bring them to the table and invite them to the table in ways that mean they will lead a conversation, not just be dragged through it.

Q585 Julian Knight: Do you think they actually understand that they have skin in the game in that respect?

Frank Sesno: Some. They have all gone to extraordinary lengths to hire new people and to take a harder look at their social responsibility in a journalistic and public information way. I have heard from many of my friends, colleagues and people I know who have gone to work there from the journalism community. They say, “Wow. This is just a completely different culture. This is a culture of engineers and technocrats, and the idea of larger editorial, social concerns is a sort of foreign language in many cases.” That is more than a minor point.

Q586 Julian Knight: I open this to any panellists who want to answer. The first amendment was mentioned, but is there not a bit of a dichotomy here in the fact that, obviously, there is no first amendment in other countries? These are American companies acting on a global scale that have their impact. From a British and probably a European perspective, we effectively see large companies that produce large profits in the United States and employ large numbers of people in the United States that are having an impact on the very fabric of our society—potentially, in the last few years, in a deeply negative way, particularly when you have Russia on the borders of Europe. Do you think that if there is to be any form of regulation—or any form of stepping up the requirement to give things such as consumer rights on data in Europe—companies such as Facebook, Google and Twitter will have to up their game globally because they still want to trade in those countries? This takes it away from the



American debate in that respect. It means that they would actually have to do something. Any thoughts?

Amy Mitchell: Absolutely. One of the biggest challenges is the degree to which anything that is going to happen in the digital space is global at this point. In the world, we have very different ways of viewing and of applying or not applying law to it. At a legal session—one of these small conferences at the Abrams centre at Yale earlier—all the folks from a legal perspective said that the greatest challenge over the next decade would be figuring out how we are going to have any kind of global approach to this that can cut across boundaries. But that has to exist. Not only the content but the audiences are constantly going cross-platform, cross-country and across the seas. In the US, one of the greatest tensions is between a very long-standing, solid support for freedom of expression and the freedom of the press. If we compare that to 38 different countries that we survey, in many of these areas of freedom of expression, the US sits far outside the global median.

Now Americans are feeling the tension of that up against what is misinformation. Two thirds of US adults agree, even across political lines, that made-up news is causing a great deal of confusion about the basic facts around current events. That is where the tension exists now. Whether the solution is regulation or something else, nine in 10 US adults are now getting some of their news online, which means those who are more or less connected to the news on a regular basis, more or less aware of what is actually going on in current events, more or less digitally savvy in terms of how to make their way through the information they see and more or less politically driven and politically motivated.

Julian Knight: Anything else?

Frank Sesno: I want to comment briefly on something you said.

Julian Knight: It was quite a long question.

Frank Sesno: You talked about the damage that this causes in societies. We need to be very careful here. Obviously, what is bad information or subversive information in one place is vital information in another. These remarkable platforms provide vital information, not just in our politics but in health, medicine and all kinds of things. We need to factor that in and define these terms, which is vital.

Q587 **Julian Knight:** That is a really good point. Do you think therefore that we need to talk less about the user experience and the consumer experience, which is what every platform talks about, and talk more about consumer rights and your right of ownership as an individual—a freedom right? Is that the way in which we can square this circle?

Amy Mitchell: Any perspective is going to raise a lot of challenges to what gets included in certain types of information or not. Claire talked about the spectrum of different definitions. To put what would be completely made-up news in perspective, we did a case study recently of stories about immigration that people linked to on Twitter, and 42% came

from what we around this table would call identified news organisations, most of which were legacy, not just digital natives. Very few were identified with any outlet that was named on a fake news list. It is the same if we look at the coverage of established news organisations of the Trump presidency during the first 100 days. We saw very different assessments of the actions of the presidency based on the audience make-up of that particular outlet. There are a lot of different kinds of misinformation, and whether a user or consumer is going to think about their right to one type of content versus another type of content can get very complicated.

David Carroll: The consumer rights aspect is really significant, and we saw it play out with the Russian investigation here in the US and Facebook and Twitter's response to the so-called information operation. It came from, in some ways, citizens applying pressure to the companies to disclose whether they were exposed to foreign propaganda. The citizens pressured the lawmakers, who then pressured the companies in the hearings, to ask, "Would you tell American citizens if they engaged with foreign propaganda that was impersonating their own citizens?" Facebook and Twitter have taken some steps to tell people whether they have been exposed to that. That begins to move us in the direction of whether we are debating censorship or privacy rights, disclosure rights and the right to know who is behind things. The same applies to who is paying for advertising and the transparency of advertising. Are we pushing the platforms to adopt the same kind of "Know your customer" principles that are required in the finance industry to prevent things like money laundering? If we have similar high-stakes disclosure issues on both the business side and the consumer side, it calls for further demands for transparency.

Q588 **Julian Knight:** That is an interesting point. It brings me to another question. What do the advertisers think of this debate right now, and what sort of pressures can they bring to bear on these companies to ensure that when they are paying the money over, they are getting what they have paid for?

David Carroll: Sure. The advertisers have a significant economic incentive to be sold accurate metrics of the audiences that they are buying, so there is a tremendous business pressure on the companies to have accurate audience measurements, which is another force to weed out bad actors in the system who are defrauding the advertising ecology by impersonating clicks—fake clicks—or amplifying fake impressions to cheat the whole industry out of its own revenue. There is a significant incentive there. The question is, why isn't that sufficient to eradicate the huge amount of fake accounts already? Facebook says it removes about a million accounts a day, and Twitter is constantly removing accounts. The anti-fraud incentives are not even sufficient to eliminate this problem, so obviously we need to do more.

Q589 **Julian Knight:** Yes, but say I wanted to set up a Facebook or Twitter account and use that account to disseminate, send out, or retweet a picture—I mentioned this in the first hearing—of a premier league goal, a



HOUSE OF COMMONS

soccer goal, thereby breaking copyright. That would come down in minutes and I would probably be banned, my account deleted and so on. Why is it that I could decide to put an absolutely horrendous lie about another person in the public domain and get another account going in order basically to retweet it and try to get it out there as much as possible, and that could just stay there ad infinitum? How is that possibly fair?

Claire Wardle: The copyright issue is huge. It was a great question to ask, and they did give a useful answer, which is that we have the technology to deal with that. I actually shared with some people the idea that we should create a database of known inaccurate content. Juniper Downs said we needed that digital footprint so that we could compare against it, and that we should be starting to create the databases to give them that.

Of course, there is a financial reason why they have moved much more quickly on what you refer to. As an individual, we are not going to sue, and sue as quickly, as the Premier League is going to sue. This goes back to the issue of advertising. The financial aspect of this question is one of the main motivations of people making this kind of content. If you talked to brands, if they were here, they would say this is one of their biggest challenges right now—the fear of having their quality brand advertised against one of these terrible sites. We are now in February 2018. The fact that we have not seen significant shifts on this, that most brands still say, “We don’t trust that we’re not going to be advertised against something that is poor quality,” is kind of astonishing. You would believe that that financial incentive had put on more pressure to stop that, but I have not seen it. We have tried to understand what is actually going to move the needle, and if finance cannot in this country, what will?

Q590 **Julian Knight:** Why is the finance not doing it? Are they not losing any money—

Claire Wardle: It is a hugely complex problem, and then there is the scale—the amount of stuff that is uploaded every single day. This is the whack-a-mole element. I work with a researcher who the other day showed me a list of new URLs that had been created this week and that are all over Facebook. The fact is that we cannot stop this stuff and every single day new domains get started. It is on a scale that is eye-watering. So yes, it is frustrating to have to say, “Why are we still here?”, but the issue is mostly the scale of this challenge and not having the computational ability to deal with it.

Q591 **Giles Watling:** It is worth remembering, to a certain extent, where all this started. This is all new stuff. These are guys who 25 years ago were sitting in sheds, coming up with these incredible platforms that we now accept as everyday. There was the big kind of Mel Gibson cry of “Freedom!” when they arrived. And suddenly we have things like the Arab spring happening, and it gave every individual a voice, which cannot be a bad thing. Now I don’t know about you guys, but I sat here today and watched people who were blinking in the sunlight—they were not



HOUSE OF COMMONS

expecting to be where they were today and this is not where they came from. My impression is that the very idea of editorial control and the five basic tenets of journalism—truth and accuracy, independence, fairness and impartiality, humanity and accountability—were not in their minds when they were creating. They are media; they are not journalists, but we are expecting to impose a journalistic regime on them.

First, I would like to get your impressions. Do you concur with me on that—on how they performed today? Secondly, how do we regulate that? Where would you come from to start it? They are not expecting it, as far as I saw.

Amy Mitchell: It is certainly the case that none of the social media platforms were begun with news in mind. None of them started as news platforms; they were all social platforms. Even Twitter was about sharing things, and conversing with your friends. So news is something that found its way to each of these platforms as people spent more and more time there and as the companies wanted to keep them there. Part of what people wanted over the course of the day was the news—to find out what was going on. Things were definitely not initially structured around being a news provider per se.

Frank Sesno: It is interesting. I was the bureau chief at CNN when AOL merged with Time Warner, and we had these very visionary conversations—"Imagine if people could just get the weather any time they wanted"—and then we started seeing some instant messaging and the speed of that. As it was when CNN was created, and when AOL was the first platform, people had stars in their eyes, but no idea of the critical mass and the power and impact of the sheer volume of correspondence.

Married to that, nowhere in this process was a journalistic mindset or set of principles imposed on the creation of that ecosystem. Journalism done the right way has gatekeepers, and those gatekeepers open and shut the gates before information goes out, not after the information is out. That is when it is done right. There is a system of accountability, a finite number of people, an org chart, and all that puts an order to it. There is no order to this social media process.

I have always pushed back against those who talk about crowdsourcing and citizen journalism. Citizen reporters, maybe—pointing a camera at something and saying, "Here's what happened," is one thing—but the true journalistic training and mindset is not something that just grows randomly. That is part of the gigantic disconnect we now have. We do not have a system of checks, balances, accountability and gatekeepers—that is the culture shock that I talked about a moment ago.

Q592 **Giles Watling:** Absolutely. So what is the way forward? It occurs to me, from discussions we have been having this week, that these companies should be linking with the academic world more, which would give them control and guidelines. How would you guide them?

Claire Wardle: I completely agree with what you are saying. The people who started these companies really believed that technology was going to



HOUSE OF COMMONS

make the world a better place. I have often said that if they spent more time down the pub with journalists, they would have realised that the world is a very dark and messy place, and being a journalist requires you to make very tough decisions every day about what you will publish, what you will not and what the impacts of those publishing decisions are.

We have talked a bit about this in terms of threat modelling. A lot of these companies have not actually sat in a dark room together and said, "What is the worst case scenario here?" I do not think that they are sharing data enough. I do not think that they are thinking through what this might be like. I think there is an element of bringing in academics, but we did not really talk today about how much of the same content just travels across those platforms. They are sitting in their silos. We have to say that they should be sharing data with one another about how the same content is travelling across.

We need them to get darker. We need them to think about what is the worst that could happen. We need them to think through what might happen and how we can respond in real time, rather than having a bunch of inquiries two years later. We should be ahead of the curve. We should be thinking about what is happening with closed messaging apps—we have not talked about those today—virtual reality and augmented reality. If we regulate today for what happened two years ago, we are in trouble. We should be looking to the future.

Q593 Giles Watling: Finally, from the evidence you saw today, do you think that they will become entrenched or are they up for change?

Claire Wardle: I think they are up for change. I do not think they know how to do this because they are terrified of opening up. They have lawyers internally who are terrified about them opening up. They work together around terrorism and extremist content, so there are frameworks that they can work together on, but they are not quite sure how to take the step from where they are to where they need to be. That is a pretty huge step, and my fear is that regulation will get in front of that. I wish we could do it the other way round.

Giles Watling: On that journey from the shed, we have arrived here.

Paul Farrelly: I want to explore briefly what steps from a consumer protection point of view might reasonably be taken to get the social media companies to open up as to how they are able to target people. I would welcome thoughts from the panel on opening up the box, as it were.

Amy Mitchell: I am not sure that I have exactly the right answer for that, based on our research, but one step that news organisations have taken is on transparency. They understand that being more transparent with their readers, audiences and users, will increase their likelihood of gaining trust and respect and having them come back and create a relationship. The other conversation happening a lot is about trust—the loss of trust that news organisations feel and so on. In many of those conversations, the steps that those organisations are taking are about transparency and

sharing more about what they know and do not know. That does not speak to the same kind of transparency we are talking about here, but it goes towards creating a relationship and a sense of trust with their user.

David Carroll: My experience with my own data has shown me at first hand the importance of the British and European data protection model and the idea of a legal subject and a legal controller, which forms the basis of creating transparency. We heard this morning that, for example, Facebook was quick to acknowledge that it has to abide by the UK Data Protection Act. As I mentioned, it is going to adopt the EU GDPR. These models show how consumer rights can be expressed through data protection rules and how those rules apply transnationally. I was able to take advantage of British law because my data was processed there, and the requirements of GDPR on all EU citizens will force companies to abide by that. But even within that, there is the significance of being able to understand what your data that you get means, how it shapes your experience and how it can be an understandable piece of data in an understandable interface.

The algorithmic accountability problem is really well connected to the data transparency problem—how is my newsfeed being shaped by my behaviour? How is my behaviour on other websites affecting the things I am seeing? I don't think that consumers have a strong understanding that all the websites they visit have their Facebook identity attached, which allows people to re-target them on Facebook. We saw that, for example, the Internet Research Agency probably used that technique to re-target Americans across platforms. General consumers do not have a clear understanding of how their data is used in various ways, sometimes against them.

Amy Mitchell: It is worth pointing out that all of this places more responsibility on the consumer. Not all members of the population are perhaps as proactive as David in looking at their information and figuring out what is going on. In terms of the privacy settings people are allowed to set and being able to turn on or off certain things, what portion of the population are actually going to take advantage of those things or take the time to understand the data about themselves and how it would be used?

Frank Sesno: I am fond of saying that we are facing an information situation not unlike the food situation we face. There is an amazing amount of food available to you, and it is all labelled, but people eat what they want to eat. We have an obesity problem in the world that is getting really bad because people eat junk food. We have brain food and we have junk food, and if we consume junk food too much, there are very serious consequences. Consumers have to be engaged and educated in a much more profound way, and that needs to be built in. Social and traditional media companies need to acknowledge that and develop a series of guidelines. Many news organisations have that. They have their ethics codes and various practices that spell out specifically from a consumer perspective what the end result should be and that lay out levels of accountability along the way. They have played out, because they have



HOUSE OF COMMONS

lived it, what happens when really bad things happen, and that has not yet happened in social media. We are on the verge with media technologies—I am sure you have seen them—where I can take one of your soundbites and change your words and put out a piece of video and it will look like you were saying something that you never said. And now we create a whole new reality, and that goes beyond just social media companies.

So this is starting, not ending, and the first thing we need to do is create partnerships, maybe imposed partnerships with some of these companies, to think about these things in a much more detailed way, coupled with the research that Pew and others are doing so that that can be brought in to these corporate cultures.

Claire Wardle: One of the best things to do is to say to somebody, “Buy an ad on Facebook,” because when they have to go through the process and see how you can target those ads, you see people’s eyes pop. As a literacy campaign to get people to understand how their data is being used, we need to come up with those sorts of ways to get people to understand what has happened.

Q594 **Paul Farrelly:** I can imagine that the companies’ default mindset might be, “We don’t charge for this. We give it away free. If you don’t like it, just butt out.” But that is not good enough.

Claire Wardle: That’s absolutely true, but when people are not paying, because they are giving up data, they do not understand that. So given that choice that you don’t pay, because you are giving up your data, if people understood that properly, they could make an actual choice about whether they want to be there and what the long-term implications of these companies having our personal data are. Maybe in the first year it is not damaging, but once you start getting older, the amount of data that they have on you starts to be able to be put together in ways that could potentially be very damaging.

Q595 **Rebecca Pow:** I want to touch on—perhaps Mr Sesno might talk about this with his media background—how much the misinformation spread by social media sites is having a knock-on effect on conventional broadcast media. It is almost teaching people not to trust anybody, so where do you go for your truth and what do real bona fide journalists do to combat this? This was raised with us by Mark Thompson at *The New York Times*.

Frank Sesno: This is perhaps the most concerning thing of all. We have created now a culture of doubt around any information, wherever it comes from in the public, and it is fed by politicians who point the finger and scream, “fake news”—again, we need to be very disciplined about that—and who refer to journalism as an enemy of the people. Not only is that not helpful, but it is very damaging and it has created both an echo chamber and a copycat effect, I am afraid, in some areas of traditional media.

We need to be careful when we talk about traditional media. The word “media” is a very plural word, and there are very many and distinct differences. In the United States, I would make a distinction between talk



HOUSE OF COMMONS

TV and talk radio, where the volume is particularly loud. Where these kinds of false reports drive an opinion-focused discussion, it has the effect of confusing vast numbers of viewers and listeners.

As to the distinction between opinion and information and what is correct and what is not, we have seen a plummet and a further bifurcation along ideological lines of trust in media just in the last year, and it is getting worse. The real concern that we all should have, and that you should have as public officials, is how you reach and inform the public and how we make sure that they have the basic information they need to participate in a democratic process where they are called upon to decide things and to elect you and make those decisions based on information and fact and not just propaganda.

Q596 Rebecca Pow: On the converse, I would also say, because this has happened to me: a Twitterstorm was caused over something I said in Parliament—it was misrepresented and put out there and my local press picked that up because it got so many hits. They didn't investigate what actually happened. They didn't necessarily check up the actual facts of the story but they thought it was a story because of the hits it had had on social media, and you could say, "Where does this end?" That is fake news that they are picking up and they are reiterating it.

Frank Sesno: It is very hard, especially in the world of real-time news, whether it is radio, cable television or some of these online things, to ignore something that is happening and playing out in real time in front of you, even if it is wrong—right? Then you have, "Well, it's generating this much traffic," or "It's out there; we have got to correct the record." So you have a very distorted kind of view. So we need to perfect what I like to refer to as "the language of live" to a very different place, where we are telling people very explicitly what that tweet was, what the controversy is around it. That is an increasingly difficult thing to do in the rush to be first and fast and loud.

Q597 Rebecca Pow: Can I ask Claire Wardle: something that concerns me slightly is that we have to be careful—as you did refer to earlier—in the legislature or regulation system not to cause another problem, because would it not be right that there will be other players waiting to come into the space, wanting to set up, wanting to play fake news? That must be a danger as well, across the world.

Claire Wardle: I think I was coming back to something that Julian Knight said. Whatever we do in Europe we can do it for the best of intentions. What it means, it will be a blueprint for all sorts of other parts of the world where there are not protections around free speech or free media—so we have to be very careful that we can be very clear about our definitions of what we are trying to protect; but it can act in a very different way in another context, so we have to be very careful about that.

To your point, too, about the media, I think one question we are not talking about or not thinking through enough is that if you are a disinformation agent, the thing that you really want is amplification by the



HOUSE OF COMMONS

mainstream media. They are targeting journalists; they are targeting mainstream media, because for them, they see debunks as a form of engagement. They are planting false information on sites—and have been since 2014 and before—that they will then resuscitate now and say, “Well, it’s been around since 2014.” Ultimately, the point of what many disinformation agents are trying to do is simply to cause confusion. It is not about any one particular rumour. It is about getting any country to a point where the citizenry say, “We can’t trust anybody any more,” and my fear is that we are getting there pretty quickly.

Frank Sesno: We are getting there pretty quickly.

Claire Wardle: I don’t know how we come back from that. We get focused on how we fact-check a particular claim or rumour. It is not really about that. It is the much wider ecosystem.

Q598 **Christian Matheson:** A very quick couple of questions, if I may. Mr Carroll, when you got your subject access request return, was the data accurate? Was the information accurate?

David Carroll: The voter registration information was perfectly accurate. The information about the election returns in my district was mostly complete but not fully complete. Then the political model was subjective. So, do I think it is accurate? There are aspects of it that are very accurate, and disturbingly accurate, and there are aspects that are impossible to understand, so part of the legal challenge in the ICO complaint is, it is important to be able to understand, for example, how this political model was generated, so that I can understand, then, how it might have been used to target me for messages.

Q599 **Christian Matheson:** All four of you might choose to answer this: when this data is being collected, because of our online activity—so this isn’t just the voter registration data, for example—is there an assumption that it will be a fair reflection of who we are and what we do, because it is being generated by our actions online, and that makes it more accurate?

David Carroll: The company says that it uses our commercial behaviour data to link it to our voter file and then make these political models. So that is why it becomes so important to understand the sourcing, because if it is the websites we visit, the products we buy, the television shows we watch, et cetera, that is then used to determine our likelihood to participate in an election, the issues that we care about most. People don’t understand that their commercial behaviour is affecting their political life.

Chair: Thank you very much.

Q600 **Ian C. Lucas:** The issue for me is how Cambridge Analytica got that information, because Facebook told me this morning that they wouldn’t share that information with anybody. From these platforms, how is it possible to extract the information if you don’t have access to Facebook accounts, or Twitter accounts, or whatever?

David Carroll: It is my understanding that the company could purchase commercial data from commercial data brokers and from ad-tech

companies and from media ratings agencies, and then use algorithms to re-identify and reconnect that commercial data back to voter profiles. That's a thing that we are trying to ascertain through the ICO and through legal challenge.

More relevant to that, we know that the researchers at the University of Cambridge who developed some of the modelling methodology and techniques did use Facebook applications to gather data from users who signed up for their application, and that is how they collected Facebook likes and used Facebook likes to predict personality, and also things like political affiliation, gender, sexuality, whether their parents got divorced, whether they smoke, drink, use drugs, et cetera. And so we have some understanding of how these techniques are used, but we are looking for more transparency to figure out how it was actually used.

Chair: Thank you. That concludes this panel. Thank you all very much for your evidence; it has been very informative. We look forward now to welcoming the final panel.

Examination of witnesses

Witnesses: David Chavern, President and CEO, News Media Alliance; Major Garrett, Chief White House Correspondent, CBS News; Tony Maddox, Executive VP and MD, CNN International; and Kinsey Wilson, Special Advisor to the President/CEO of the New York Times.

Q601 Chair: I'd like to welcome the witnesses and our guests to the final panel for today's evidence session. The final stage today is really to discuss some of the issues we have looked at today with the tech companies and with the academic experts as well, to discuss those with people representing different aspects of the news media industry.

If I could start off by asking the members of panel if they could give their view on an issue that has been giving us cause for concern as we've gone through this inquiry, and that is that public news consumption is increasingly moving out of a curated news space, whether a news bulletin, or an edited newspaper, or a news website. People are consuming news more in bite-size pieces that they discover and share via social media, and increasing numbers of people get their news via social media. The latest Pew research is that about two thirds of people in this country get news from social media sites.

So what cause for concern does that give you about news from your organisations and also about the consumer's point of view and the public's point of view about whether they are getting a fair, balanced picture from the news that they consume, if they are largely just consuming a variety of different articles that are shared with them or they are exposed to via their news feeds on social media sites? Kinsey Wilson, perhaps you could start us off and give us a perspective from your news organisation.



Kinsey Wilson: Sure. On the one hand, this is simply the reality of how people consume news and information. They are awash 24 hours a day—all their waking hours, certainly—in small and large bits of news. As news organisations, we simply have to recognise that that is how people have become accustomed to consuming news. One of the things that I think is often hard to understand inside a news organisation is that people do not come to what we do with the same kind of intention they did in the past. I think we probably had naïve notions as to how completely they absorbed and read the material we wrote and how many articles they read. In fact, now we understand that it may be sufficient to read a sentence or two sometimes—you're on the move. There are other times when people will go on their phone and read a 2,000 or 3,000-word story. It is highly variable.

The Times has declared that we have to remain a destination for our readers, so while we use these platforms as a way to expose people to the kind of journalism we do and attract newer audiences and so forth, we have been very resolute in saying we ultimately need to bring our most engaged readers back to *The Times* itself. Other organisations find it difficult to do that and consequentially are much more dependent on how these platforms display their news and information. They have effectively become the equivalent of the front page for most individuals, so while traditional news organisations are still composing individual stories, they are arranging and figuring out who sees what in what order and in what kind of time frame.

To the earlier conversations, it makes it very difficult to determine whether they are in the publishing space or merely a purveyor of content that other people have produced. Clearly, through their algorithms they are applying a level of judgment to what people see, and that has a societal impact, quite apart from the misinformation that is flowing through those platforms. How that gets regulated, what the right approach is, is difficult in part because there isn't a clear bright line between what constitutes news and other types of information. Traditional news organisations in this country are not licensed and, apart from libel and defamation laws and so forth, there is no particular regulation that applies to us. So there is a spectrum of information that has now gone into the mix on these platforms and it becomes very difficult to separate it out. I think in the quiet counsels of their offices they would just as soon be out of the news business, if they could be, and not have this kind of scrutiny.

You asked what scares them. In addition to regulation and possible trust action, it is also the prospect that what they have built is to some extent getting out of their control. Their reputation, if you look at how they score, generally Google and Facebook enjoy pretty good reputations, but that is taking a hit. There are all kinds of unintended consequences that, as engineers, they are having difficulty getting their hands around.

Q602 **Chair:** David Chavern, we have heard from Kinsey Wilson. *The New York Times* is a major global news brand, which probably gives it certain advantages in making the transition away from being a printed



HOUSE OF COMMONS

newspaper into a digital media business. Looking at the breadth of members you represent with the News Media Alliance, do you think we are going to see the consolidation of the news media business continue apace, with fewer titles and some parts of the country where they will have very limited local news altogether?

David Chavern: I think there are several challenges. Certainly, there will be, and will continue to be, a consolidation in the news business as people's attention, and the brands to which they attach that attention, are necessarily somewhat limited. That being said, there is a real fear. We talk about fake news as a national story—the President has talked about fake news, and we talk about various issues on a national scale with fake news—but I think the future of fake news will be almost entirely a local phenomenon. People have curiosity about, and interest in, their communities, but there will be insufficient reliable sources locally, and that will be filled by crazy bloggers and conspiracy theorists who have ideas about what is happening in the local council or on the local school board.

Looking to the future, one of the primary challenges, certainly for local news in the US, is, as we move into digital consumption of news to a greater degree all the time, do the local news organisations have the resources to build and optimise and connect with their audiences digitally, as compared with national players such as *The New York Times* that have great resources and great capacity to make this transition? Will local news organisations have that same capacity? There is a lot of concern about that.

On the local news front, we could say, "These people are very close to and attached to their communities, so that should be an advantage," and it is an advantage, but the question for the future will be one of capacity: will they be able to make the investments to make this digital jump?

There is also the question of what will be their relationship with their distributors. In the past, when they had a printed news product, they created it and literally handed it to their audience; it was the ultimate direct relationship. This is now, to use a tech term, disintermediated; these are sources that put out content and there is somebody—notably Google and Facebook—who stand between them and their readers.

Kinsey is absolutely right that, particularly in the US, the news business has not been regulated; it is that great First Amendment we have. But I often say that Google and Facebook are our regulators. Their rules about distribution and delivery impact our businesses every day. It will be a continuing challenge for local news organisations in particular to figure out how to work their way through disintermediate delivery—for want of a better term—and to have the capacity to make this technological jump.

Q603 **Chair:** Turning to Tony Maddox and Major Garrett, from a broadcasting point of view, two things follow on from this that have been of interest to us in the various meetings we have had since being in America this last week. First, do you feel that the key issue will be one of trust, and with things like augmented reality giving people the technical capability to



HOUSE OF COMMONS

create not just fake stories, but fake videos and films, there is an opportunity for the traditional broadcasters to be the more trusted gatekeepers of properly checked out and sourced information, because there will be a lot more, not just written stories distributed digitally, but faked images?

Also, do you feel that the broadcasters need to do more to innovate their services to reach audiences online? We had a very interesting meeting in New York with NowThis in the media, a Facebook-based news channel with 2 billion viewers a month. Has it just got lucky as an early adopter, or is that model difficult to replicate for a more traditional broadcaster?

Major Garrett: I think everyone on this panel would agree that the dividing lines that used to exist—and that existed when I first entered this business in 1984 as a local newspaper reporter in Amarillo, Texas—between newspapers and television have blurred entirely. CBSnews.com is just as important a news source as our broadcast programmes “CBS This Morning” and the “CBS Evening News”. We have enormous click rates all through our CBS.com content all day and all night.

Every journalistic outfit, large or small, has tried to find a way to merge video, audio and publishing together, and to do it credibly. In response to some of the earlier questions, I have been in this business since 1984, and when I worked in a community newspaper, it was the community talking to itself about its values and sense of perspective, and there was one choice in that community to find your news: the newspaper, and then the local television stations. The choices are now all over the place. We exist in a world where consumers are—as I heard Frank Sesno say a few moments ago—eating junk food. They do not think it is junk food; they think it’s spinach. They’re making a conscious decision to say, “No, what I was being fed”—maybe from CBS or NBC or CNN or *The New York Times*—“was the junk food and what I am finding alternatively is my good food.”

We have to acknowledge that that separation from our journalistic organisations that was probably built on a foundation of having limited choices has gone away, and we have to now compete over credibility and trust. That is an hour-by-hour, day-by-day pursuit.

There is nothing that CBS News, NBC or ABC can do about someone who mashes together video in a dishonest way, except broadcast that which is legitimate. There is a slippery slope here—maybe that is not the right term, but there is something that media organisations have to take account for. When there is a hurricane, tornado or snowstorm, we will take content off Facebook and Twitter that people have shot themselves, and we will put it on the news. Why? Because we weren’t there, and they were. We get it for free. We may put a teeny microscopic line in there that credits the source, but we are taking content and turning it into journalism at no cost. We don’t have to pay for a crew, an assignment editor or any of the loaded costs we used to pay for to create our own wholly owned broadcast product.



HOUSE OF COMMONS

We are part of this dynamic, and it behoves everyone in our industry to understand and own up to some of that. We, in certain ways, are compromised commercially, but there are certain ways in which we profit—or, if not profit, gain—from all this content that we then transmute into journalism. I have a slogan I use all the time on two of these points. One is: content is content, journalism is journalism; they are not interchangeable concepts. The other is on news consumers. Whenever I give speeches on this and people ask me about what to believe and what not to believe, I always leave them with this thought: if from the time you wake up to the time you lay your head down to go to sleep, everything you read and consume in a news context leaves you pleased, happy and verified, you're doing it wrong.

Tony Maddox: I would echo that. Probably the greatest single concern I have, with all the big pluses that have been identified with the increase in the availability of news, the sources of news and the variety of news, is that it has enabled people to create their own little media bubble in which they only really get material that is consistent with their beliefs. That encourages fake news, frankly, because within that sphere, anything that is consistent with it gets thrown in there—"We don't like this guy, so I've read this thing or heard this terrible thing." My concern is that you don't get to read provocative things. You don't get to read things that you want to throw across the room because you disagree with them. You need to be exposed to things other than what you and all your friends think.

In terms of news and politics—I see this in the UK, even from the US, with things like Brexit—it has almost become like supporting a team. You are pretty much in one camp or the other, and people are accused of not being a true supporter or devoted enough because they can see some aspect of the other person's argument. How has it got to this? How did we get to a situation where you have to sign up for 25 different things because you broadly think this or that? Most normal human beings do not do that. It is really unhealthy if we create a news consumption environment where everything that comes in is consistent with a particular argument. Of all of the concerns I have, that is a primary one.

You made an important point at the start of this, which is that faking stuff to look like really good stuff gets easier all the time. The technology is much cheaper. The programming is much cheaper. It would be very easy for someone with pretty rudimentary skills to make something, and your average layperson is not going to know whether ABC, CBS or CNN did it or some person in their room did it.

Major Garrett: Can I interject with a thought? It is a gloomy topic, but I want to identify one thing that has happened in our country. The proximity of people with their own cell phones has revolutionised understanding of a certain issue in this country and fundamentally changed news organisations' relationship to that issue, and that is police violence in this country.

I was a police reporter for six years at three newspapers in our country earlier in my career. If there was a police-involved shooting, and I wasn't



HOUSE OF COMMONS

there, I was dependent upon the orientation to the facts, circumstances, timing and evidence provided by the source most responsible, but also most potentially liable—the local police department. The presence of citizens with their own independent means by which to video what happened has transformed not only that conversation but the relationship between what happened and what others who are not—possibly—defensive about what happened say happened. It has fundamentally altered, I think in a very positive way, not just the community's understanding of what happened but that of journalists who are trying to figure out what happened. There are a lot of things that we have identified that are problematic with all this, but that is one area I would like to interject has changed and changed for the better.

Q604 Chair: To follow up, Major Garrett, there is another side to that too, which is the pressure being placed on news organisations to get news out fast because citizen news is breaking all the time. We have certainly seen examples in our country where people have got it wrong because of that speed pressure. Is that something you recognise?

Major Garrett: Oh, absolutely. I can give you a classic example. After the Boston marathon bombing, a tremendous amount of information was provided by citizen journalists or those who were monitoring police scanners, not only right after the bombing itself but on that very fateful Friday night when there was a successful effort to find and apprehend the suspects. Things were appearing on social media based on what somebody heard on a police scanner. We were very cautious at CBS and a lot of other news organisations were cautious, but it was out there, and when it is out there, there is this deeply philosophical and highly pressurised conversation about what we do with that. Are we close to verifying it? What is our standard of verification in this very heated environment? I kept counselling my own network internally that I have listened to police scanners—I did that for six years. Whatever you hear from one cop on the beat is relevant, but it is also a very thin understanding of what is going on in the totality. We have got to be careful. The circumstances create enormous pressure, much greater pressure than I ever experienced in the earlier part of my career in broadcasting, which I started at CNN in 2000. Part of this continuum between what is believable and what is hyper-pressurised is the ability to say philosophically, "I am prepared as a news organisation to be second and right. I am never prepared to be first and wrong."

Q605 Julie Elliott: I am particularly interested in the interference with elections and referenda of disinformation campaigns that are both organised and non-organised. The non-organised bit is very much what you were talking about, Tony, where people reinforce their own beliefs and they retweet or repost and reinforce.

If we could talk about the organised first, you have all been in journalism a long time. What do you think has really changed in the last couple of years? Does it feel worse? Does the interference feel worse? Or is it that we simply know that the interference is there and it was probably always there? What do you think?



HOUSE OF COMMONS

David Chavern: Let me start with a reaction. One of the good things about the digital consumption of news is that you can consume a lot. When I was growing up, we had the paper on the driveway and the half-hour news and that was the news. That was all your window into the world. Now you have access to so much and that is an absolute good thing. But also, we all have crazy relatives that over the course of dinner have weird conspiracy theories and lies—man didn't go to the moon and whatever the theory is. But you knew that information was different from what landed on the driveway and was on the television. It came from a different place. One was professionals and one was not professionals.

One of the digital challenges is that it is all put into a blender and fed to you in very similar ways. It puts a tremendous onus on the readers—the consumers—to differentiate. There are insufficient indicators to readers about what is good stuff and what is not good stuff. We can talk about that. It provides a tremendous opportunity for people who want to manipulate the public. They can use these big digital pipes to feed out garbage news. By the way, the most powerful kind of fake news is not the kind that says that aliens have come down. It is something that is somewhat off—something that people wish were true and that feeds their biases. People can feed you a bunch of information and, unless you are a really careful reader, you do not understand that it is coming from a bad place. I think that it has opened up an opportunity for people to manipulate the public much more than in the past. But there are things that we can do about that, such as indicating what good journalism is and what it isn't. I will pause and let other folks respond.

Kinsey Wilson: I think, from my perspective, the biggest concern is targeting and the lack of transparency around targeting. That is probably even more problematic than the spectrum and sheer volume of information out there. We have quite a good understanding on the commercial side of our business of how powerful those tools are. We spend a lot of money on Facebook. It is the most efficient way for us to acquire subscribers, simply because of its enormous wealth of data and ability to target. It can be used in a variety of ways.

In these conversations, it is important to distinguish between what is showing up in people's news feeds and in search—which is very much in the control of these platforms as they write their algorithms, and try to tune the difference between social connection and authority or relevance and authority—versus the ability of those who pay to target particular audiences on these platforms. You have got into the conversation about potential solutions. One important place to start to unpack this issue is certainly at the level of transparency, individuals having control over their own data, and perhaps data portability. Without that, so long as these platforms are in complete proprietary control of that information, there is a risk that regulation will have a lot of unintended consequences, as well as good effects.

Tony Maddox: To your point, it has been a fascinating day listening to all the key issues. I do not envy you trying to work your way through all



HOUSE OF COMMONS

these different issues. From where I am sitting, I respectfully suggest that nothing is more fundamental than the sanctity of elections, and we are talking about the idea that third parties could interfere in elections and that you then do not have full and fair elections.

Having spent more than a dozen years at the BBC before I moved to CNN, I have quite a bit of experience of the UK and UK elections. From outside, there are a couple of things that I always point out when I explain the UK elections to my American colleagues. One is the Representation of the People Act. There is nothing like that formal in the US. For all its flaws, it has enabled a broadly fair template that all broadcasters and major news organisations need to subscribe to. But the other thing that I point out is quite technical—that, because of the number of constituencies, there are a lot of small majorities in the UK. Look at the number of constituencies in the UK with a majority of less than 1,000 or maybe 1,500 votes. Those constituencies are enormously vulnerable.

If people are following the election primarily via social media, and people who have a desire for an outcome one way or another decide to target that area, it is very easy—as you have heard—to get geo-specific information so that you can ensure that this story hits all your users in this particular constituency. It is very easy to get a profile of people's interests. If it's an immigration story, for example, you can ensure that everybody who might have an interest in immigration in this constituency gets hit with this wave of stories. There has to be some kind of defence against that. I think the media companies themselves are faced with an enormous challenge. I say "media companies", but they call themselves "tech companies", I notice; whatever they may be, they do have a big challenge here—they have grown very, very quickly—and I am sympathetic to that challenge. We know that they have great geo-targeting technology, and some of this is just going to be a simple question of employing people.

On this idea that we cannot tell whether a story is a legitimate news story, well, you can, if you hire some local journalists—a lot of local newspaper journalists have been laid off—for an election and say, "Just check this story out. Is it from *The Guardian*? Is it someone's opinion?", which is all fine, "Or is it not like that? In which case, we should take it down." It's almost worth making a priority of the elections, certainly in the US. The idea that Russia interfered with the election is a profound issue. For all the other chatter and stuff that goes on, you keep coming back to the idea that everybody is seriously disconcerted that that may have happened. That is where I would start, if you like, trying to get my arms around this.

Major Garrett: Two quick observations: one, means and motive. I think that actors have long had motives to see if they can disturb or influence the American political process. The means is different. I have a lot of conversations with political actors now, because I spend a lot of time on Capitol Hill and at the White House, and we have started having this conversation. I try to remind them that we have given them the means by which to fulfil their motives. By being so archly partisan against one



HOUSE OF COMMONS

another, we have created, I believe, in our dialogue in the last 10 years in this very country, a much wider opportunity for fake news or propaganda, or things that divide us still more, to have traction. Our own political discourse has created an opening that I do not believe existed before. I think one of the reckonings to which our own political system has to come is the way we talk about ourselves and the way we speak about differences of opinion.

That is point one. Point two: we have to at least acknowledge the possibility that there will be behavioural adaptations to all the things we are talking about. I use my three children as an example: 22, 21 and 17. All three, in varying degrees, have grown up in this telephone, digital world, and they are already demonstrating signs of exhaustion—psychologically and otherwise—and I believe that there is a very real possibility that as those who grew up with this adapt differently they will begin to send signals to Google, Facebook and others about what they believe is credible and not credible. That possibility—of them changing, quite separately from any regulatory regime anyone might ever create—gives me some sense of optimism that we are all collectively, especially the younger generation, who have experienced this and know no other reality, coming to a reckoning about the underlying basis of shared news, that basis being all the third-party verifiers you trust, and that that trust could be reanalysed and maybe redirected in other ways.

I am optimistic about that. Perhaps that is unfounded optimism, but I see it in my own children and the behaviour they express, not only for themselves, but on behalf of their friends. They are looking at this world—at the tech companies and this whole process of how the world comes to them—and the novelty has worn off, and I believe that there is a great possibility of adaptation, at the consumer level, for all these things we now look at and assume will always be the same. If we have learned anything in the last 10 or 15 years, it is that things will not always be the same.

Q606 Julie Elliott: Thank you. A number of you have mentioned fact checking. One thing I am concerned about is that, with this fast-paced 24-hour news thing, often on social media things are put out there that aren't true. It was mentioned this morning, in respect of the Boston bombers, that information was out there that wasn't true—that people had been arrested, but they hadn't. We have had similar things in our country over the last 12 to 18 months. One thing that concerns me about what I call traditional journalism is that, with fewer people working on newspapers, particularly local newspapers, and fewer people working in the big broadcasting companies, and with the urgency to get a story out there and up, even some of the bigger concerns are not always checking—they'll take something and think, "That's all right. We'll get ahead of the game on this." What impact do you think that's having, and do you think it's something that the more traditional media outlets perhaps should be looking at again—going back to some of the basics to make sure things are correct before they put them out there?

Tony Maddox: I can see this changing now. I have worked at CNN for 20 years and when I first worked at CNN one of the things we always used to

do in a big breaking news situation was, if we were the first to break the news, we would get a quick PR campaign out, saying "CNN at 11.38 said this". We don't do that any more. There are a couple of reasons why we don't do it any more. One is primarily because stories don't break first on TV, unless it is your own original investigation, in which case the time it goes out is irrelevant, because it is an original piece of journalism. In terms of breaking events, natural disasters, crimes or whatever it might be that breaks first, it will break via social media. It will break via Twitter or something like DataMinr, which is a Twitter/social-media-based news tip-off service.

Julie Elliott: I have never even heard of that.

Major Garrett: We have!

Tony Maddox: That's the way it goes. They will be here next year.

The point is, the primacy of saying "We got it 10 seconds before everyone else" is going away, and I think you will see news organisations who—it is also worth bearing in mind that in a traditional news organisation, when you get something wrong, and it does happen, it is awful. People think, "Oh, they got it wrong, it is fake news" and everything else. When a mistake is made it is terrible. Everybody feels dreadful about it. There are full investigations. There are real consequences. Sometimes people lose their jobs. It is a really bad, bad thing to happen. So people will do a lot, a great deal, to avoid that.

I think as more and more sources of instant news become available people evolve in what they look for from established brands. For all of the frenzy that surrounded fake news, and all that has gone on in the US, what is worth bearing in mind is that traditional news organisations like CNN, like *The New York Times*, have enjoyed improved fortunes. The amount of consumption of our products is on the rise and it is significant that people put together their own portfolio of news. They might see something on social media, but sometimes folks are quite savvy and they think, "Well, I have seen that there, but it would be different again if I read it in *The New York Times* or on 'Anderson Cooper' than something I just read on Twitter". So I do think that the traditional media has a role, but that role is to be increasingly the source of trusted news. That is inconsistent with trying to get everything out instantly, or operating in an under-resourced way. There is no future in that.

Kinsey Wilson: I think we have adapted in many cases. One of the conventions we use is we are very quick to tell people both what we know and what we don't know. Because these kinds of breaking events have such currency in social media you can't go dark until you confirm everything, necessarily, because people look to you to try and understand what is actually happening and what they should believe and should not. So we have all adapted and understood. I think television, probably, was much better at it than newspapers, because they operated in a live environment for years. They have had to adjust to the changes in social



HOUSE OF COMMONS

media. I think we have all learned pretty well how to act responsibly. It is tougher in organisations that are very thinly resourced, certainly.

David Chavern: A couple of interesting trends. First of all, I think there is a view that—just the absolute value of traffic. There has been this idea of more traffic: better. That was based on digital advertising models where “We get more traffic, we can expose more advertising, we get more money from advertising.” Google and Facebook now command the lion’s share, and growing, of all digital ad dollars, so you are finding news organisations thinking about their future in terms very much of a subscription, direct economic relationship with the reader. I think one of the enriching parts to that is really caring about long-term trust, and not wanting to get these things wrong; so I do think the impetus to get it now, so we can grab the traffic, is dissipating to some degree. One of the unfortunate parts is Google and Facebook are getting all the digital ad dollars, but there is this renewed emphasis, I think, on subscription and direct relationship with readers.

On the fact-checking side, journalists are the fact checkers. There are fact-checking organisations that do good things, but, at the end of the day, the best fact checkers are paid professional journalists. I still think that there are insufficient indicators, too, when people are consuming news online as to what is coming from professional paid journalists and what is coming from something else. I certainly always encourage Google and Facebook in this regard. We need to get much better at giving people indicators about what is coming from professional news organisations and what is coming from something else.

Major Garrett: To that point, in newsrooms five or six years ago, it was not at all uncommon to have a large flat-screen television with click rates for every one story from that day—if you were in a newspaper. I was at a magazine at the time. Back in that time—which seems not so long ago, but in terms of adaptation it was a while ago—stories would be updated, and those updates would carry corrections. For a while, there was a kind of—I think—debasing informality between a correction and an update. This phenomenon of having clicks theoretically translated to ad dollars: as the industry has become more accustomed to consumer habits, it has learned that it needs a relationship—it needs longer views—and not just clicks in and out of stories in order to monetise that. So updates now are much more transparently disclosed as corrections if that is in fact what they are. News organisations have become much more aggressive about distinguishing between the two, and much more transparent about that, in order to re-establish that idea of a relationship. So it is not just clicks any more that dominate what is or is not a successful story. What you need is a lot of credible stories that create a long-term relationship and a longer-term view of whatever your product is, because that is actually, in a very old-fashioned sense, how you monetise your journalism.

Julie Elliott: Finally, Major Garrett, in your written evidence, your last but one sentence says “credible journalism will always outlast incredible politicians.” But what we are dealing with here is incredible journalism.

Major Garrett: With all due respect—

Q607 **Julie Elliott:** We don't take it personally. What we are dealing with is incredible journalism—so not credible journalism—impacting on our political process. What are your thoughts on that?

Major Garrett: Well, that it has found commercial traction. Information that a segment of our society—your society—finds either reaffirming or valuable has found a marketplace, has found a business model that will pursue it. That is a reality. I can only take—not only professionally, but philosophically—the long view on that, that that which is not only true today, but true tomorrow and a year from now, will ever live in starker contrast to that which is not, and those who are responsible for that which was true will gain perhaps incrementally but steadily over time. If I didn't believe that, I could not do what I am doing, and I certainly couldn't do it in the atmosphere that I am doing it in now.

Q608 **Brendan O'Hara:** A couple of questions. Over the past week, we have heard quite a lot about the financial effect on the traditional mainstream media of the loss of circulation and advertising revenue due to the sharp rise in social media. Please feel free to comment on that, but I also ask you for your thoughts and reflections on the long-term effect on the social and cultural wellbeing of the United States—where you are—of this loss of influence of the mainstream media in the public sphere. Do you have any concerns?

Kinsey Wilson: I think we should have profound concerns. The business model for journalism has been disrupted for the past 20-22 years, and essentially, with the ability of anybody to publish and the barriers to entry into the publishing business having fallen, advertisers are able to go direct to audience and not need media as an intermediary. These are larger phenomena that go beyond this current conversation or even Facebook and Google, as dominant as they have become.

It is very hard to unwind that. It has led—David is probably in a stronger position to talk about this—to what we talk about in the US as news deserts, where entire communities are largely devoid of the conventional coverage of public institutions, courts and so forth that you could count on in the past. It is not clear, as much as they are trying, that either established news organisations or the many start-ups—for profit and not-for profit—that have been funded over the years have as yet found a credible, sustainable model for that kind of local journalism. We are finding models at a global—certainly international—scale, and it is clear that those that are succeeding generally are succeeding on the strength of trust and connection to their readers, reviewers and listeners and by the quality of what they produce, but there are issues with scale in local communities that leave real questions as to how the fundamental journalism will be funded. How you maintain it and inform society is essential to democracy.

David Chavern: I have a couple of comments. First, one of the central problems that we are wrestling with here, and that certainly Google,



HOUSE OF COMMONS

Facebook and Twitter are wrestling with—Major Garrett referred to it in his comments—is that news is a different kind of content, even if you don't want to use the word "content". It is important. What I mean is that a disappointing television show or cat video is just that: it is disappointing. Disappointing, as in low-quality or fake, news can destroy civil society on a broad scale and certainly at a community level. It is existential to our civil society that we have access to quality news. I don't think that is something that the tech giants understood when they got into this world. I think they are trying to figure it out and have not done that yet. We are all wrestling with the fact that there is no such thing as a free news business, and we need quality news, or else we will not have a civil society.

When you look at the local level, to take what I referred to before as the idea that we may well move to subscription models, we are losing out on the advertising side to Google and Facebook. Print advertising is declining. We will go to subscription models. There are some fantastic examples of that, including what the gentleman right here, Kinsey Wilson, and *The New York Times* have been able to do with print digital. Digital subscription in particular is extraordinary. But you have to wonder, at the community level, what is the capacity of a community to pay subscriptions that will support real journalism in that community.

If you do not have real journalism in a community, you quickly get to fake-news world locally. I can give an example. I was on the city council in my little town of Falls Church, Virginia, for several years. We had a journalist who covered things in the city council. If he was not there, I assure you, there would have been a crazy blogger in the back room who insisted that the school board consisted of primarily aliens or something, and they would get that information out. There will be people feeding curiosity about the community. But if it is not actual professional reporters, it will be fake news, which will have disastrous impacts on those communities. There was a great piece in *The New York Times* a few weeks ago about a fake news phenomenon in Twin Falls, Idaho, that was brutally destructive to the local community environment there.

So I do not think there is a known answer right now for local and community news. We are struggling to help find the answer, but we are going to have to find it, because news is important. If we lose it, we will lose quite a bit.

Major Garrett: I commend to the Committee's attention two examples in America. They are small examples, but they possibly point to the future. One is the *Texas Tribune*, a foundation newspaper supported on a non-profit basis. It has been in existence for quite some time. It predates a lot of the phenomena we are talking about, but there is an example that has cropped up in the last year and a half in Las Vegas, Nevada. I was an employee many years ago at the *Las Vegas Review-Journal*. A journalistic website entirely, the *Nevada Independent* has built itself entirely on non-profit donations from the community and from the businesses within that community as an independent source of journalism. It has no offices, no printing press and no delivery trucks. It has journalists and laptops and a



HOUSE OF COMMONS

website, and it is beginning to work. That is one possible model for filling the gap at the local news level.

My career started as a newspaper man in three different cities in America—Amarillo, Las Vegas and Houston. All of them had two newspapers competing against each other when I was there, and they have one struggling newspaper now. What our country and anyone affected by these titanic shifts in the marketplace have to comprehend is that journalism is a public trust and a public service and that it has a public value. If the public values it, the esteemed citizens of said community have to figure out a way to support it, and there are small examples of that beginning to emerge in our country that give me some optimism.

Q609 **Brendan O'Hara:** Was that the *Nevada Independent*?

Major Garrett: Yes.

Q610 **Brendan O'Hara:** Okay, thank you. I have one final question. An intentional consequence of this fake news phenomenon has been the loss of trust in any, sometimes all, sources of news—the loss of the traditional gatekeeper, if you like. It really is one of the most worrying aspects of everything that we have come across. With the greatest respect, gentlemen, you were never trained in being able to combat the loss of all trust and this rise in fake news. What are you or your organisations doing? How do you begin to rebuild that trust, which would appear to be diminishing by the day?

Tony Maddox: I don't know that it is. There is always the balancing act of appearing arrogant and aloof and saying that these things aren't having an impact, and then overreacting and saying that the sky is falling. I don't think that the sky is falling. I think there has been a systematic, focused attempt to de-legitimise traditional media, because the stories that traditional media are reporting are unhelpful, so if they can be de-legitimised, you reduce the impact of those stories. There is a concerted, focused, ongoing attempt to do that, and the way to respond is to double down on what it is that you're good at.

You could make a very strong case that the past year and a half or two years have actually been something of a golden era in American journalism. If you look at the number of major stories that have been broken from traditional news outlets—I certainly include *The New York Times*, *The Washington Post* and organisations like my own—one story after another has been fundamental to the American people, the American Government and the way in which America works, so I am not going to give in to the idea that we have somehow been de-legitimised.

Certainly, as a company ourselves, having been consistently attacked, we do our own individual research. It is extensive and repetitive—four times a year—to see what the impact has been on audiences. I have to tell you that the information we have is that the impact has been next to none, and that is certainly supported by stuff like ratings, revenues and everything else.



HOUSE OF COMMONS

We would be foolish in the extreme not to recognise that these constant attacks do cut into the core of what we do, but we cannot panic for that. We shouldn't change what we are doing; we should double down and reinforce the values of what we do. The evidence is that that is working for us.

Kinsey Wilson: I would certainly agree with that. If you dig into the polling, there is no question that trust in media has been declining if you ask broadly about the media as a category. If you begin to ask about particular publications that people consume or local publications and so forth, you get different signals. I saw a study yesterday from Edelman that actually indicated that they were seeing a bit of an upswing in the trust in media and a decline in the trust in the platform. A lot of it is driven circumstantially by what is going on and what is in the news.

We have absolutely doubled down on doing what we do best, which is digging deep and being as thorough and as comprehensive as we can. Despite the economic pressures on us, we have added to the newsroom in Washington and beyond, and we are seeing the results in terms of the number of people who are subscribing, the traffic that we get to the *Times* and even, frankly, being able to sustain our print circulation in an environment where there is secular decline.

There are two stories going on simultaneously. The things that we have been talking about do erode people's trust in a whole variety of information sources, but there is also an appetite for something they can latch on to, use to cut through the noise and actually trust.

Q611 **Brendan O'Hara:** How much of that, though, is to do with polarisation in communities, where each person runs to his or her preferred flag? Is that what you are seeing?

Kinsey Wilson: There is certainly a fair amount of political polarisation, to be sure. I do not know that people are simply gravitating to their preferred political choice. It is complex. Certainly you can find cohorts of people who are doing that and who are attached to one political outlet, but there are also people who are reading more broadly and have an appetite to really understand issues across the board.

Major Garrett: It feels to me a tremendous opportunity. I really mean that. For the last 20 years or so, this idea that there is something untrustworthy about the mainstream media has gone unanswered. We now have a lively conversation in our country about what journalism is and what it isn't. Your presence here is indicative of a global interest in the foundational set of principles around credible journalism and trustworthy reportage. I tell my colleagues all the time that this is a very emotional time—it feels very heavy—but it is also the best opportunity we have ever had in our careers to do everything that Tony and Kinsey just talked about: to lay it out and show what it is every single day. The audience has never been more interested in what it actually is that we do and how we go about doing it. There are market pressures and partisan pressures, but there are enormous opportunities to do this right and show the way.



HOUSE OF COMMONS

Chair: Thank you. I am conscious that we are about to enter our seventh hour of evidence. I know a few other colleagues want to come in, but we will try to be reasonably efficient in the rest of the time we have.

Q612 **Paul Farrelly:** I will try to be brief. Tony's point about where you start—elections—was well made. In my constituency, Newcastle-under-Lyme, which is—small world—Tony's birthplace, I have a magnificent majority of 30. I am one of 11 members of what I less than affectionately call the "under-100 club". If you sway a few hundred votes, you change the Government in the UK. That is how tight it is.

I was an investigative journalist, so I believe that a fair, accurate, free and responsible press is absolutely vital. Of course, the new digital media organisations perform a very useful role. They would not be so popular if they did not. But as far as your companies are concerned, they are certainly grabbing your audience and taking large chunks of your revenue. They are arguably going to grow in influence as traditional outlets struggle, but as we have heard today, they take very little responsibility for some of the sadder aspects and absolutely no liability. Do you think that is right? Do you think that is fair? Do you think it is fair competition?

David Chavern: First of all, it is a reality. These are amazing products that give our news brands access to many, many people, and they have built amazing technologies. People always say, "Are they publishers? Are they media companies?" I would apply any and all of those labels to them. I primarily call them attention companies: they want to access as much of the public's attention as possible. They want more of it tomorrow than they had today. They do own responsibility for what they are using to get your attention.

Now, do I want a world in which Google and Facebook are editors and publishers and express editorial perspectives on what is on their platforms? No. But I think they can do a much better job of helping the user to separate the wheat from the chaff and understand what information is coming from what sources, with indicators of credibility.

So they can do much more to help resolve the "fake news" phenomenon. I think the difficulty is that as engineering and technology companies, they like engineering and technology solutions. This may be an area where they have to do something crazy, like hire people. But do they have some responsibility for what their users are exposed to when they are using those users to make money? Yes, they do.

Kinsey Wilson: If the question is whether the competition is fair or not, it is important to distinguish between the platforms. I would include Apple and Amazon in this as well. There are those that behave more like walled gardens, Apple and Facebook in particular, where they are trying to keep their users in their environment and maximise the time spent there. Others, like Google and Twitter, are redirecting traffic back out. In the latter instance, the economics of media are tremendously complicated, but we at least have an opportunity to build our business if they are effectively acting as distribution platforms and

common carriers, if you will, and not simply as a walled garden. The argument we have made to them, and to Facebook in particular, is: either build a set of APIs and a system that allows us to manage our business on your platform, or pay us for our content. One or the other. Give us the tools to allow us to build a business there or treat it as a closed cable system and actually pay us for the content. Their size is such that we don't have a huge amount of leverage—even *The New York Times* doesn't have an enormous amount of leverage to make a lot of progress on that front—but lumping them together is not necessarily always helpful, particularly when trying to understand the economics of the ecosystem.

Paul Farrelly: Tony, you can read the question whichever way you like.

Tony Maddox: Just a few thoughts really. I agree with David that they are going to have to hire some people. You can't solve everything with a logarithm and they are going to have to realise that. I also think there has been a positive. CNN, in many ways, was one of the original disruptors. No one had done a news model like CNN. Once CNN did that everything in the news business changed and this is like the next iteration of that. There is an argument that we were maybe due for a shake-up. There has been much soul-searching inside the news business both in the UK and the US. Why did we miss Brexit? Why did we miss Trump? If we spent more time looking at some of the conversations taking place outside traditional media we might have been more alive to what was going on there, so I think there is a salutary lesson in that.

Then I think there is a general principle. I can't really be too specific on this, but no matter what business you're in it can't be a principle that we don't care about the truth. Whatever it is you do, the truth has got to be important. You can't say, "If it's a lie, we can live with it." I do not know how you get round that, but that can't be a basis of the business.

Major Garrett: A couple of real quick answers. I mentioned I was in the newspaper business earlier in my career. Back then, most newspapers you worked for, even small ones, clocked about 12% to 11% profit every year. Why? Because they sold classified advertising. Classified advertising was an embedded financial part of the business model that never went away. People say, "Did the internet kill newspapering?" Well, newspapering isn't dead, but it wounded it. But not the internet. Craigslist and eBay wounded American newspapering, because they gave everyone who used to go down to their local newspaper and pay \$15 for three lines of agate type to sell their geraniums, puppies or golf clubs an alternative place to send that money. That took 40% to 35% out of every newspaper in America within the course of about two years. I ask you: what business model can sustain 30% to 40% loss of revenue in two years, find no means by which to replace it, and maintain the same standards, size and scope? The answer is zero. That is a business reality—a competitive reality—that none of us can do anything about, but it is real. That created smaller newsrooms, fewer ambitions and more openings, because the conversation in communities around the country was becoming less robust.

As was mentioned earlier, when there is no coverage of your local community, everything masquerades as a partisan issue at the national level. I guarantee

you that there is no Republican or Democrat national orientation to whether or not a bond is let in your community to fix the water pipes, but if there is no local journalist there to tell you about that, it feels like it. That loss of revenue, ambition and scope at the local level was filled in by some of the phenomena we are talking about.

The second point I made earlier was about adaptation. Google, Facebook and other large companies have for a long time created a false impression about their relationship with their user. My children have already figured this out. The relationship with me is not really about me; it is about you—what you learn about me and then what you sell about me. The great future for journalism is to say, “Our relationship with you is completely different. You subscribe to us and we’ll tell you about your world,” whether that is a balloon festival in the springtime in your local community or a civil war in Syria that is in its seventh year. “We are not going to mine you for the rest of your life and sell everything you’ve ever purchased, every thought you ever had and every piece of music you ever downloaded to 17 other companies. We’ll tell you about your world.” The sooner we get to that place and that relationship in this adapted world, the stronger we will be, the better off we will be and the more we will be able to push back against all these other fake news phenomena.

Q613 Jo Stevens: Tony, I was smiling at what you said about how lies can be the basis of a business. That is effectively what we heard from Twitter in their evidence this morning. They are quite happy to have lies on their platform and for their business to be run and sustained on the basis of that. My question is to David. You mentioned indicators to help users understand what has come from professional journalism and what has not. Trying to be positive and looking ahead at potential solutions, one of the things suggested to us in the meetings we have had this week relates to *Consumer Reports*, which is an independent rating agency you have here. It is a bit like our *Which?* magazine in the UK. Do you think that an independent rating agency for information and news outlets that measures facts against fiction and reporting against opinion, with some ownership identity information—the equivalent of a nutritional label on a packet of food—would work? If outlets had poor ratings, that might influence them to change their behaviour.

David Chavern: I would start with something simpler and then get to that. The simpler thing would deal with what, for want of a better term, is called brand suppression, whereby the brands that are attached to news stories get minimised, and frankly, with the scraping and theft of data, the brands get very confusing to users who are not working really hard at figuring things out. The platforms could easily work on making brands more prominent and having more brand equity with the users.

Then you get to whether there should be a capacity to determine the wheat from the chaff in terms of professional news organisations. I think there should. While it is often described as an impossible task, I don’t think it is. First, there is already line-drawing in the journalism business. There are journalists who get credentials to do some things, and organisations that can’t get journalists credentials to do other things. You could have pretty basic and encompassing standards that relate to, for



HOUSE OF COMMONS

example, whether an organisation actually hires and pays journalists. What is their ability to respond, and do they respond to comments and corrections?

A number of different things could be developed that would not be very difficult to apply, particularly for companies with the resources of Google and Facebook, to determine professional news organisations from non-professional ones. Again, we could be pretty encompassing. Frankly, just paying journalists is a big one.

Q614 Jo Stevens: But without that element of independence, can we trust those platforms to do that, based on their behaviour?

David Chavern: There are various decisions about this. There is a group called the Trust Project and there are other folks looking at indie-share quality. There are a lot of ways in which we could get to that, but I would then give extra credit in the algorithm to folks who are professional journalistic organisations. That would not mean getting rid of or censoring the garbage, but it would give more opportunity for readers to get reliable journalism.

Kinsey Wilson: Google is actively working with us and with other news organisations to try and attach signals that essentially give a sense of authority to a story, as well as looking at the relevance of the actual information that somebody queried, and they are willing to open-source it and make it more widely available to other platforms. Whether that works in a social environment that is really about the connection between individuals and less to do with surfacing a particular piece of information is an open question.

Q615 Rebecca Pow: These are only very short points. As a former journalist and broadcaster, I am heartened that you are so optimistic that there is a future for journalism and good reporting, although on the converse the whole reason why a lot of it has gone downhill is because there simply is not the advertising to put into the newspapers and provide the backing, because the advertising has all been drawn on to social media, as we have been hearing all week. Some 20% of worldwide advertising is now on social media. How are you going to tackle it, just briefly Mr Maddox?

Tony Maddox: Me alone? It is an easy one to finish on! We have seen the ability of social media companies to provide advertisers with incredible detail to allow them to target in a way we have never seen before. That is a major game-changer.

Q616 Rebecca Pow: I meant to add of course that the advertising has gone on to social media, and the advertisers want to advertise there because so many people are looking at it, and so many people are looking, as we have heard, on the sites that get the most hits, which tend to be—but not always—the least reliable and those that have the most potential to be fake.

Tony Maddox: This business is changing very quickly. One of the counterintuitive things with digital media is that lots and lots of hits can be



empty calories as far as advertisers are concerned. This is a quick question to finish, but it is actually the beginning of an entirely new session. My company at the moment—Time Warner, which owns CNN—is in the process of trying to merge with AT&T and is going through court cases and everything else. One of the primary reasons why these two big companies are trying to do that is because they are trying to combine the data that one has and the content the other has to put together an offering that is along the same lines as the social media companies. At the highest levels of these companies, they realise they are going to have to try to provide a similar kind of metrics. So this is completely the opposite end from the local TV question; this is about how corporate America responds to this major shift in the plates of the advertising industry, and how that will work is that massive companies will be prepared to merge and work together and reinvent themselves in order to respond to that shift within the business. That is quite a big, pompous answer, but that is really where the answer starts and it will break down from there.

Q617 Rebecca Pow: And my suggestion, to sum up, which perhaps we could put to the gentleman from *The New York Times* is: surely the answer, which you've alluded to, is that all you journalists link up with Facebook, Twitter and Google and provide this high-powered, high-profile journalistic content for the engineers?

Kinsey Wilson: On the advertising side there are efforts under a trade organisation called Digital Content Next to put together something called TrustX, which basically aggregates quality publishers into an advertising exchange that commands higher cost rates per thousand than you typically get from programmatic advertising and things like that. At the end of the day, it is about scale, whether it is broadcast organisations combining with mobile carriers and so forth, or other news organisations trying to pull together to create an ad market of their own. That provides perhaps a partial answer. We did our earnings call today and announced that we now have \$1 billion in subscription revenue, on a company that is \$1.5 billion, so two thirds of our revenue is coming through subscriptions. Ten years ago, that was probably flipped, or not even—it was more like 80:20. So at least in certain quarters that is part of the solution as well. There have to be alternative sources of revenue beyond—

Major Garrett: That is the relationship I was talking about. That is the relationship model for the future of journalism.

David Chavern: I'd make two points. First, digital advertising would be a whole other seven hours-plus of hearings. If you think the news business has been disrupted, you should talk to people on the buy side of advertising and see how their world is different from what it was 15 or 20 years ago.

The piece where you can build optimism about news and journalism is that the audience is bigger than ever. There are more people consuming more news than ever, period. They do it because they can. You have access to so much more. That is something to build around.



HOUSE OF COMMONS

Ultimately, there is a lot of work and thought being put into subscription models. That is a big part of why we have been asking Google and Facebook to facilitate subscription through their platforms. They are experimenting with that but have not yet delivered on it. That could go a long way, instead of getting in the way and disaggregating the relationship with the reader, to helping us solidify it.

People do want news. They do value it. There are some positive things, with *The New York Times* being a primary example. We have to all express a lot of concern about the future for local and community news. The relationship with the reader will be there. The capacity to pay is going to be an open question for a while.

Q618 Chair: Thank you very much. I very much agree with what you said, Tony, about the models of the established media businesses changing to respond to the power of the tech platforms we have been talking about today. It is probably a matter for a separate inquiry, let alone a separate evidence session, but it is certainly something of interest to the Committee. The interest that the Committee has taken in things like the Disney-Fox-Sky deal also reflects the same phenomenon.

Major Garrett: Mr Chairman, forgive me. I am a working journalist. I have a deadline. I need to ask your forgiveness—I have to leave to get back to the White House. Thank you so much for your time, and forgive me.

Q619 Ian C. Lucas: I am going to be very brief. I was very interested in what you said, Tony, about the merger. One of the things that really strikes me about this is the monopoly situation that we effectively have, with Facebook and Google having 60% of US digital ad spend and 20% of total global ad spend at the present time. So we have a monopoly situation. Someone earlier on briefly mentioned anti-trust. How has this monopoly situation been allowed to develop?

Tony Maddox: I'm going to pass on that. I mean, we are just about to have a court case on it.

Ian C. Lucas: Okay, that's fair.

David Chavern: I think it's all about data. The question is, if you are in the data business, does that lead to a potential natural monopoly, meaning if you have a little bit more data than the other guy, does that give you huge advantages in attracting ad dollars? Is there a natural monopoly there? I think there are a lot of academics doing research on that. There are a lot of analyses as to whether traditional views of anti-trust, which are usually, particularly in the US, consumer benefit views, apply in data-centric markets. I don't know where that will ultimately end up, but it is true that when your advantage is data, if you have a little bit more than the other person, you have a lot more, and that is always going to be true.

Kinsey Wilson: I think there is something in the question of data portability. Without portability, you tend to get into a situation of lock-in,



HOUSE OF COMMONS

where the biggest established players simply can't be dislodged and there is no room for competitors to gain a foothold. If I can't take my information and move it to a better service, and I instead have to wait for all my friends and everybody else gradually to adopt that service, the chance that it is going to take hold and present some sort of competition to these largest players is minimised. I quickly get out of my depth as to whether these things are better dealt with through anti-trust or not. There are differences between European and American views of anti-trust and so forth. The data issue—not simply the command of it, but who owns it and whether it is portable—is at the heart of a lot of this.

Q620 Christian Matheson: I have just a brief question, and it is more of a general one. To an extent, you have addressed it. Do you see that the overall desire for news and information has grown? David, I think you said you felt it had. Is there an element of optimism that the overall desire for quality journalism, as opposed to other forms of news, is growing? Can you demarcate any of that by demographics within society?

David Chavern: We have a foundation called the American Press Institute that has some data on news consumption demographics. It is clear that there used to be tropes—that people don't consume news anymore or that young people don't consume news. You don't hear any of that now because it is so patently ridiculous. Part of that is because of what is happening in the world and people's focus on national news stories, the US presidency and other things. But there is wide recognition that people will always remain curious about what is happening in their world and their communities.

There is a lot of data showing that because so much news and information is available, it is consumed a lot. Now, it is consumed differently—by millennials, for example—than how we consumed it as young people, right? It may come in a Facebook feed, with beach selfie, cat video and Syrian civil war all in your news feed. At all ages, people consume massive amounts of news, which is a fundamentally good thing. It is something that you can potentially build a future on. This is not a business where we are lacking customers. This is a business in which how we made money from creating quality journalism has been incredibly disrupted and broken, and we have to rebuild what the future looks like from the ground up.

Tony Maddox: I agree with that. If you are going to spend your life in the news business, you have to be optimistic. This job will wear you out otherwise. I think that there are a lot of grounds for optimism in terms of news consumption and in terms of people entering the business, including the quality of people coming into the news business, the number of people who want to get jobs, the variety of jobs and the diversity of people trying to get into the news business. Those are all enormously encouraging trends.

This has re-emphasised to big brands the importance of what their brand can be, and of investing in quality journalism. That has resonated. It is good for all kinds of reasons, but it has also proved to be a good business model, as you have heard. A lot of good things are happening right now.



HOUSE OF COMMONS

News consumption is higher than it has ever been. If I look out here, people think of the *New York Times* as a newspaper and CNN as a 24-hour TV news service. Behind me, these things are viewed as digital propositions and they are pretty much platform-agnostic in terms of how they are consumed. But the key thing is that they are being consumed, and in greater levels than ever before. That is a great cause for optimism.

Kinsey Wilson: I would agree with that. There is such a bleed between news and other forms of information—the cat videos are thrown in with headlines and so forth. Richard Gingras has made the point in other settings that a lot of fact-based information that used to reside in newspapers—sports scores, weather and things that you knew to be demonstrably true—is no longer there. There is an enormous obligation on us to explain what we do, and to help people understand the difference between how different types of information are gathered and how the profession works. We have been working very hard to be much more transparent about that. We have been trying to give people a sense of how we go about our jobs and why it is important, and not to simply assume that that is well understood.

Chair: Great. Thank you very much for your evidence. That concludes our questions this afternoon.